Assessing Macroeconomic Tail Risk

Francesca Loria* Christian Matthes† Donghai Zhang‡

January 6, 2022

Abstract

Real GDP and industrial production in the US feature substantial tail risk. While this fact is well documented, several questions remain unanswered. Is this asymmetry driven by a specific structural shock? No. We show that the 10th percentile of the predictive growth distributions responds about three times more than the median to *both* monetary policy and financial shocks. What mechanism can generate this asymmetry in the data? We discuss nonlinear VAR models and a nonlinear equilibrium model that are capable of matching our empirical findings. Furthermore, we provide empirical evidence that allows us to differentiate between two competing theories.

Keywords: Macroeconomic Risk, Shocks, Local Projections

JEL Classification: C21, C53, E30, E37

We thank Marvin Nöller for excellent research assistance on a previous version of the paper. We also want to thank Danilo Cascaldi-Garcia, Dario Caldara, Fabio Canova, Todd Clark, Giovanni Favara, Domenico Giannone, Nobuhiro Kiyotaki, Ricardo Reis, Erick Sager, Rosen Valchev, and Alexander Wolman, as well as seminar/workshop participants at the Chinese University of Hong Kong, University of Bonn, International Finance Workshop on Quantile Regressions at the Board of Governors, 2019 IAAE conference, 2020 SNDE conference, "Modeling The Macroeconomy in Risky Times" Workshop in St. Louis, "Macro-at-Risk" session of the ECB Working Group on Econometric Modelling, IMF's "Monetary and Capital Markets (MCM) Policy Forum" and SEACEN for helpful comments.

Disclaimer: The views presented herein are those of the author and do not necessarily reflect those of the Federal Reserve Board, the Federal Reserve System or their staff.

*Board of Governors of the Federal Reserve System, francesca.loria@frb.gov.

†Indiana University, matthesc@iu.edu.

‡Institute for Macroeconomics and Econometrics — University of Bonn, donghai.zhang@uni-bonn.de.

1 Introduction

Recent contributions in macroeconomics and finance (Adrian, Boyarchenko and Giannone, 2019; Giglio, Kelly and Pruitt, 2016) have highlighted that growth in measures of US output are highly asymmetric and feature substantial *tail risk*. To illustrate this feature, Figure 1 plots the quantiles of the one-year ahead forecast distribution of growth in Real Gross Domestic Product (GDP) and Industrial Production (IP), commonly called "growth-at-risk". We replicate the results from Adrian, Boyarchenko and Giannone (2019) based on quarterly GDP as well as the monthly GDP measure that we will rely on in our benchmark analysis, which is based on the model of Caldara, Cascaldi-Garcia, Cuba Borda and Loria (2020).[1]

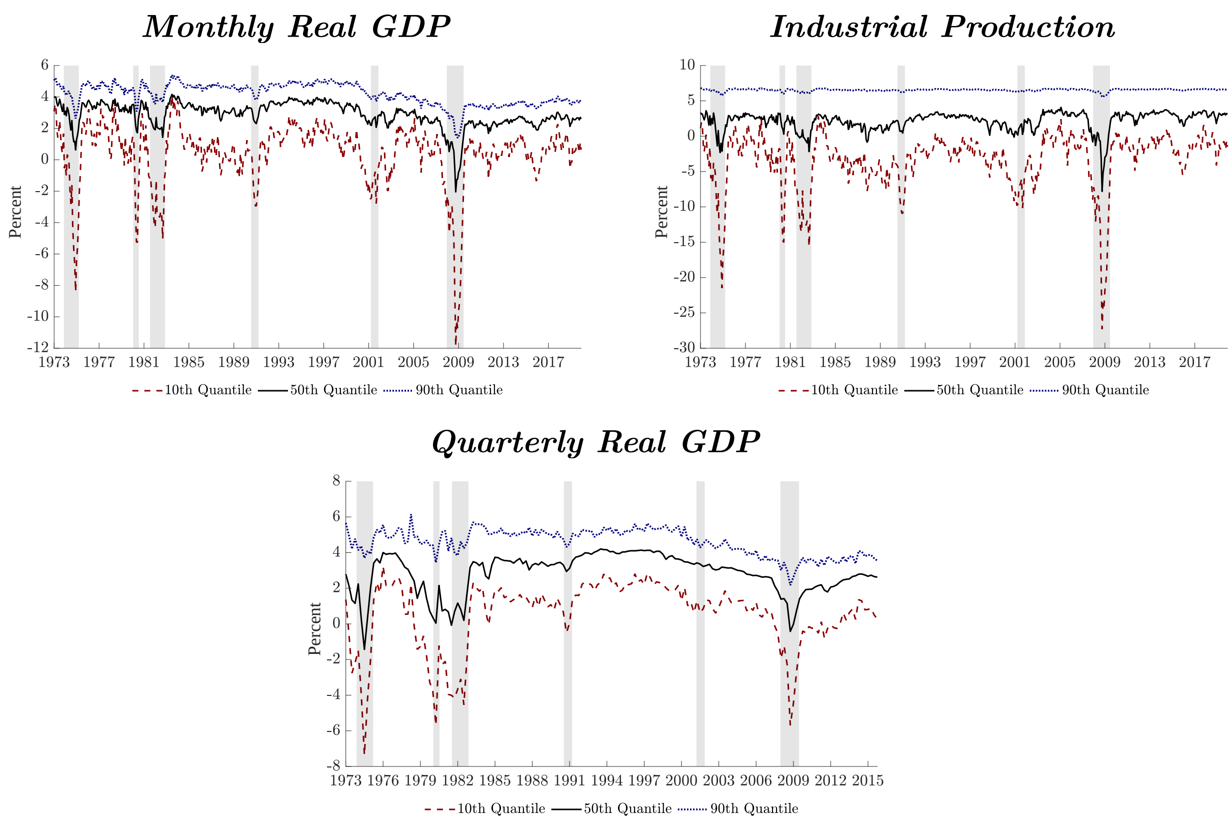

Figure 1: Quantiles of One-Year-Ahead Growth for Various Measures.

Note: Quantiles of Average Growth over the Next Year of Monthly Real GDP Growth and Industrial Production Growth as in Caldara, Cascaldi-Garcia, Cuba Borda and Loria (2020) (Top Panel) from January 1973 to December 2019, and of Quarterly Real GDP Growth as in Adrian, Boyarchenko and Giannone (2019) (Bottom Panel) from 1973-Q1 to 2015-Q4. Grey-shaded bars indicate NBER-dated recessions.

[1] Giglio, Kelly and Pruitt (2016) focus on asymmetries in industrial production growth due to systemic risk. Figure 1 shows that this asymmetry is robust to using other specifications than those in their paper.

All panels convey the same asymmetry – the 10th quantile moves substantially more, in a recession, relative to the median or the 90th percentile.[2] While this fact is well documented, several questions remain unanswered. First, is this fact driven by a specific structural shock? Second, which mechanism is behind this asymmetry? We tackle these questions in two stages.

In the first stage, we derive a convenient two-step estimation routine that merges ideas from Quantile Regression (Koenker and Bassett, 1978) to determine quantiles of the distribution of interest with Local Projection (Jordà, 2005) to assess the effects of monetary policy shocks and shocks to financial conditions on the distribution of real GDP growth. We devise a bootstrap algorithm to capture all relevant statistical uncertainty in this two-step approach and that might be of independent interest to some readers. We find that responses to *both* shocks considered here show substantial asymmetry - over the first year the 10th percentile moves on average 3 times more than the median in response to the structural shocks.

In the second stage, we devise a conceptual framework to interpret these results by means of two data-generating processes (DGPs), which we also use as laboratories to test our empirical machinery. The first DGP, a "semi-structural" threshold model, reveals crucial insights into how to discern between various mechanisms that can drive growth-at-risk. In particular, while several configurations of the model can give rise to observationally similar responses of GDP growth to structural shocks, we show that some can be discriminated by looking at the response of the conditioning set of the quantile regression. Guided by these insights, we estimate the responses of the conditioning variables of our empirical GDP growth quantiles and find that structural shocks exert nonlinear effects on GDP growth through both financial conditions and macroeconomic activity. The second DGP, the nonlinear DSGE model with financial panics by Gertler, Kiyotaki and Prestipino (2019), confirms that a common propagation mechanism amplifies downturns, thus supporting the plausibility of our empirical findings. We also use these DGPs to highlight specification choices in our framework.

Our analysis is motivated by concerns that economic policymakers and market participants have long harbored: They are generally not only worried about what changes to economic conditions will do to the economy *on average*, but also how these changes affect both the probability of *large* harmful events occurring as well as the magnitude of these events.[3] We borrow the concept of *value-at-risk* from the finance literature to operationalize the concept

[2]Details on the construction can be found in Section 2.1.

[3]For research emphasizing that the Federal Reserve is concerned by downside risk, see Kilian and Manganelli (2008). For direct evidence of a policymaker thinking about downside risk (which we will use synonymously with tail risk), see the March 2019 speech by Lael Brainard.

of macroeconomic tail risk. To be more precise, we follow Adrian, Boyarchenko and Giannone (2019) and study the distribution of macroeconomic risk by estimating a quantile forecast regression of real GDP growth over the next year for various quantiles.
Standard impulse response functions in linear models such as Vector Autoregressions (VARs) are not built to answer these questions as they track average outcomes. While fully parametric VARs that specify the probability distribution of the one-step ahead forecast error could be used to derive such measures of risk, these VARs put many restrictions on the behavior of forecast densities.[4] Our goal is to provide a flexible, yet simple framework that can directly tackle these issues. A flexible approach is useful when studying GDP growth because GDP growth is far from Gaussian, as can be seen from papers that interpret the non-normality in GDP growth as coming from fat tails and/or stochastic volatility (Fagiolo, Napoletano and Roventini, 2008, Cúrdia, Del Negro and Greenwald, 2014, and Justiniano and Primiceri, 2008). Our approach is constructed to be flexible through the following modeling choices: In the initial quantile regression stage, we model each quantile separately instead of assuming a specific distribution for the forecast distribution of real GDP growth. In the second stage, we use local projections to impose as few restrictions on the data generating process as possible.[5]

In devising a conceptual framework to study the sources of macroeconomic risk, our paper also provides key insights on how to replicate growth-at-risk patterns in theoretical models. So far, there have been limited efforts in trying to rationalize results from the growth-at-risk literature and to use them as guidance for economic modeling. One such effort is Adrian, Liang, Zabczyk and Duarte (2020) which adapt a standard New Keynesian model to include "financial vulnerability" driven by movements in endogenous and forward looking financial conditions and calibrate it to match growth-at-risk patterns. Another example is Aikman, Bluwstein and Karmakar (2021) which build a semi-structural New Keynesian model with financial frictions and obtain a fat-tailed distribution of GDP. What sets our approach apart is that we specifically build data-generating processes that explain our empirical evidence on the asymmetric effects of shocks on the growth outlook, as outlined above.

[4]Even non-linear time series models such as VARs with stochastic volatility and time-varying parameters along the lines of Cogley and Sargent (2005) and Primiceri (2005) impose substantial structure. Carriero, Clark and Marcellino (2020) propose an extension to the standard stochastic volatility structure embedded in these models that helps to capture tail risks.

[5]As shown by Plagborg-Møller and Wolf (2019), local projections and VARs asymptotically estimate the same impulse responses, but are on diametrically opposite ends of the bias-variance trade-off in finite samples.

The paper is organized as follows. Section 2 gives an overview of our methodology to estimate the quantiles and the impulse responses of those quantiles to structural shocks. We then introduce the specifications and data transformations used in our empirical analysis in Section 3, which also presents our main results. In Section 4 we lay out a conceptual framework to rationalize and shed light on the potential mechanisms behind our findings, and show more empirical results to support these mechanisms. We offer concluding remarks and suggest avenues for future research in Section 5.

2 Econometric Methodology

Our approach consists of two steps, which we outline in this section. First, we compute quantiles of the predictive distribution of one-year-ahead GDP growth. We focus on the 10th percentile and, as reference points, the median and 90th percentiles. We interpret this 10th percentile of the predictive distribution of future GDP growth as *macroeconomic risk*. Our second step then uses local projections to assess how structural shocks affect macroeconomic risk and the predictive distribution of average future real GDP growth more generally. In Appendix A we provide intuition for impulse responses of quantiles.

2.1 Conditional Quantiles

Quantile regression models (Koenker and Bassett, 1978) are a flexible tool to quantify risks surrounding the growth outlook and to study its determinants.[6] The first step in our analysis is to use this methodology to compute conditional quantiles of future GDP growth as in the quantile regression application in Caldara, Cascaldi-Garcia, Cuba Borda and Loria (2020). Their framework is inspired by Adrian, Boyarchenko and Giannone (2019), with the important difference of moving to monthly frequency so as to allow a timely assessment of developments in financial markets and the real economy. This is done by estimating GDP growth as well as a financial and a macroeconomic factor at monthly frequency using a Mixed-Frequency Dynamic Factor Model. The data are illustrated in Appendix B. We use this monthly measure of GDP as our benchmark because GDP takes a broader view of real activity compared to Industrial Production and a monthly frequency is helpful to both enlarge our data sample and minimize issues of temporal aggregation when it comes to the identification of shocks. However, our results are very similar if we switch to a quarterly fre-

[6]For an introduction to the quantile regression methodology, see Koenker (2005).

quency and use standard quarterly data on GDP or if we use Industrial Production instead of the monthly measure of GDP, as shown in Appendix E.

Let us denote by $\bar{\Delta} y_{t+1,t+12}$ the average (annualized) real GDP growth over the next 12 months and by FIN_t and MF_t the financial factor and macroeconomic factor, respectively. Formally, the conditional future GDP growth quantiles are estimated from a linear quantile regression model whose predicted value

$$q_{\tau,t} \equiv \widehat{Q}_\tau(\bar{\Delta} y_{t+1,t+12} | FIN_t, MF_t) = \hat{\alpha}_\tau + \hat{\beta}_\tau FIN_t + \hat{\gamma}_\tau MF_t, \quad \tau \in (0,1) \tag{2.1}$$

is a consistent estimator of the quantile function of $\bar{\Delta} y_{t+1,t+12}$ conditional on $\{FIN_t, MF_t\}$.[7,8] For the estimation of the parameters we use data from January 1973 to December 2019.

The time evolution of the one-year-ahead predictive GDP growth distribution can be illustrated through the fitted values for the 10th quantile (left tail), the median and the 90th quantile (right tail). As previously shown in Figure 1, these quantiles (top-left panel) are similar to those obtained by Adrian, Boyarchenko and Giannone (2019) (bottom panel) at quarterly frequency (originally run for the period from 1973:Q1 to 2015:Q4) by conditioning on the National Financial Conditions Index (NFCI) and GDP growth.

One novel finding in Caldara, Cascaldi-Garcia, Cuba Borda and Loria (2020) is that not only deteriorating financial conditions but also decelerating economic activity, once measured by the informationally rich macroeconomic factor instead of GDP growth as in Adrian, Boyarchenko and Giannone (2019), make the growth outlook more vulnerable. In this sense, this supports the evidence in Plagborg-Møller, Reichlin, Ricco and Hasenzagl (2020) that financial conditions are not the (only) determinant of downside risk.

2.2 Impulse Responses

Since the fitted quantiles summarize not only the median but, most importantly, the tails of the future GDP growth distribution, they constitute a measure of macroeconomic (downside

[7]Formally, the dependence between explanatory variables x_t and a quantile of y_t is measured by $\hat{\beta}_\tau$:

$$\hat{\beta}_\tau = \underset{\beta_\tau \in \mathbb{R}}{\operatorname{argmin}} \sum_{t=1}^{T} \left(\tau \cdot \mathbb{1}_{(y_t \geq x_t \beta)} |y_t - x_t \beta_\tau| + (1-\tau) \cdot \mathbb{1}_{(y_t < x_t \beta)} |y_t - x_t \beta_\tau| \right), \ \tau \in (0,1)$$

where $\mathbb{1}_{(\cdot)}$ denotes the indicator function, taking the value one if the condition is satisfied. Note that no distributional assumptions about the error term are required.

[8]A similar approach using factors in quantile regressions can be found in Giglio, Kelly and Pruitt (2016).

and upside) risk. Our interest lies in investigating whether and how these measures of risk respond to aggregate shocks. We estimate responses of different future GDP growth quantiles to a variety of aggregate shocks by applying the local projection method of Jordà (2005). As a baseline, we run the following linear regression:

$$q_{\tau,t+s} = \delta^s_\tau + \theta^s_\tau \; shock_t + \Psi(L)^s_\tau \; controls_t + u^s_{\tau,t+s}, \quad s = \{0, \dots, S\} \tag{2.2}$$

where $q_{\tau,t+s}$ is the τth quantile computed in the previous section, $shock_t$ is the structural shock of interest, and $\Psi(L)^s_\tau \; controls_t$ is a lag polynomial of control variables which include the lagged quantiles and model-specific controls that we will discuss in detail later.[9] Note that there are two distinct notions of "horizon" in our application. First, the horizon in the quantile regression h, which we keep fixed at one year and which captures how forward-looking our measure of risk is. The second notion of horizon is s in the local projection, which we vary as we trace out how risks respond at different horizons to a shock at time t. The response of quantile q_τ at time $t+s$ to a *shock* at time t is then given by θ^s_τ.

We construct the impulse-response functions by estimating the sequence of the θ^s_τ's in a series of univariate regressions for each horizon. Confidence bands are based on the bootstrap procedure described in Appendix C, which controls for serial correlation in the error terms and the estimation error in the quantiles.

At this point it is useful to contrast our approach with another approach that aims to combine quantile regressions with local projections, an approach advocated for by Linnemann and Winkler (2016). We interpret the 10th percentile of average future GDP growth as a measure of downside risk and we then ask how this measure of risk reacts to different shocks. We study a number of shocks and find it useful to use the same quantile (or measure of risk) for all shocks we study in our local projections. Linnemann and Winkler (2016), instead, are interested in one shock only and model the conditional quantiles *conditional on, among other things, a fiscal shock* and thus include the shock directly in the quantile regression. Linnemann and Winkler (2016) cannot distinguish between the two horizons h and s that we emphasized above (given that they ask a different question, they probably would not want to).[10] An advantage of their approach is that they do not require a separate step to

[9] To be specific, we have $\Psi(L)^s_\tau = \Psi^s_{1,\tau} L + \Psi^s_{2,\tau} L^2 + \dots + \Psi^s_{p,\tau} L^p$ so that $\Psi(L)^s_\tau \; controls_t$ only contains p lags of the control variables.

[10] Another approach in empirical macroeconomics that uses quantile regressions is introduced in Mumtaz and Surico (2015), who use quantile autoregressive models to study state dependence in the consumption-interest rate relationship. Recent work that combines quantile regressions with VAR models to estimate impulse responses is presented in Chavleishvili and Manganelli (2017) and Kim, Lee and Mizen (2019).

estimate the quantiles. An advantage of our approach is that it allows us to focus on the shocks' impact on future growth that operates solely through the conditioning variables of the quantile regression (in our case, macroeconomic activity and financial conditions). This also comes with the benefit of reducing noise and thus increasing the robustness of the empirical model. Indeed, the one-step approach requires more data to obtain reliable estimates as it tries to estimate conditional quantiles that also condition on specific values of the shocks of interest. Reassuringly, a one-step approach gives similar results, as shown in Appendix E, likely reflecting the fact that the financial and macroeconomic factor are a sufficiently rich set to describe growth-at-risk dynamics. Furthermore, we later confirm these findings in Section 4 by means of a Monte Carlo exercise.

3 The Response of Tail Risk to Macroeconomic Shocks

We estimate how quantiles of average GDP growth over the next year respond to two aggregate shocks, a monetary shock and a financial shock. We choose these shocks because they represent very different sources of economic fluctuations - if the response patterns are similar across these shocks this gives us confidence that a common propagation mechanism is at play. We identify these shocks via *standard* instruments in the literature. Next, we introduce the specifications and data transformations. More details on our data sources are provided in Appendix B. A large battery of robustness checks can be found in Appendix E.

Monetary Policy Shocks Our benchmark choice are the monetary policy shocks extracted from a version of the Gertler and Karadi (2015) proxy VAR which uses the Miranda-Agrippino and Ricco (2020) shocks (updated up to December 2019) as a proxy for the monetary disturbance in the model.[11] Details on how we extract the monetary policy shocks from the Gertler and Karadi (2015) VAR are provided in Appendix D.

We estimate the following regression for the sample from January 1986 to December 2019:

$$q_{\tau,t+s} = \delta_{\tau}^{s} + \theta_{\tau}^{s}\ shock_t + \Psi(L)_{\tau}^{s}\ shock_t + u_{\tau,t+s}^{s}, \tag{3.1}$$

where $shock_t$ denotes the monetary policy shock and twelve of its lags are controlled for. In

[11]We use Miranda-Agrippino and Ricco (2020) shocks as instruments to deal with the potential confounding problem in the high-frequency monetary surprises that may arise due to the Fed's private information, see, e.g., Nakamura and Steinsson (2018), Jarociński and Karadi (2020), Miranda-Agrippino and Ricco (2020), and Zhang (2019).

Appendix E we show that other measures of monetary policy shocks, such as the Antolín-Díaz and Rubio-Ramírez (2018) and Romer and Romer (2004) shocks, that allow to start the sample in 1973, give rise to very similar results.

Credit Spread Shocks We take the Gilchrist and Zakrajšek (2012) excess bond premium (EBP) updated by Favara, Gilchrist, Lewis and Zakrajšek (2016) to construct an aggregate credit spread shock. This shock can be identified by setting the controls in the local projection appropriately.[12] For the identification of EBP shocks, we assume that no other shocks can affect the excess bond premium on impact within a month.

We estimate the following regression for the sample from January 1973 to December 2019:

$$q_{\tau,t+s} = \delta^s_\tau + \theta^s_\tau \; EBP_t + \Psi(L)^s_\tau \; \left[\; EBP_t, \; \pi^{cpi}_t, \; \Delta y_t, \; i_t \right]' + u^s_{\tau,t+s}, \tag{3.2}$$

where we include twelve lags of all controls. In Appendix E we show that starting the sample in 1986 as the monetary policy shock, gives rise to very similar results.

Summarizing Uncertainty Our two-step approach is subject to two key sources of statistical uncertainty: not only are the coefficients in the local projections step subject to estimation uncertainty, but the quantiles that are used as data in the local projection step themselves are estimated with uncertainty. We thus design a bootstrap procedure to obtain confidence intervals that captures the uncertainty of both the quantile regression and the local projection step. The full details are provided in Appendix C.

Interquantile Ranges The uncertainty around the IRFs of single quantiles estimated in separation from each other cannot readily answer the question of whether quantile responses are statistically different from each other because they represent marginal responses that disregard the correlation between the responses of the various quantiles. This is why we run, for both shocks, another regression in which we recognize that the quantiles come from a common data generating process and thus compute the impulse responses of interquantile ranges directly. The specification reads:

$$q_{\tau_1,t+s} - q_{\tau_2,t+s} = \delta^s_\tau + \theta^s_\tau \; shock_t + \Psi(L)^s_\tau \; X_t + u^s_{\tau,t+s}, \tag{3.3}$$

[12] For a further discussion of how timing restrictions such as this can be incorporated into local projections, see Barnichon and Brownlees (2019) and Plagborg-Møller and Wolf (2019).

where τ_1 and τ_2 are the quantiles of the interquantile range and X_t the shock-specific controls.

3.1 Results

We present the impulse responses of the one-year-ahead GDP growth quantiles and interquantile ranges in Figure 2. In the top panel, we plot the impulse responses of the quantiles after a contractionary shock. To make the responses comparable across shocks, we rescale the responses across quantiles such that the median falls by 25 basis points on impact (this procedure does not distort the sign of the response). The key takeaway from Figure 2 is that there is a clear asymmetry in the response to all shocks we consider: *The 10th percentile moves more than the median, which in turn moves more than the 90th percentile.* The responses of the difference between the 10th and 50th quantiles are statistically different from zero, whereas the same cannot be said about the responses between the difference in the responses between the 50th and 90th quantiles.

These differences are economically sizeable. To give a sense of the quantities, Table 1 shows both the average ratio and the maximum ratio of the impulse responses of the 10th percentile relative to the median for the first year (in other words, we look at the value of the red line in the right panels of Figure 2 divided by the black line in the same figure). The average ratio is 2.6 for the monetary policy shock and 2.7 for the credit spread shock. The maximum ratios over the first year are similar since ratios are fairly stable. The response of the 10th percentile for both shocks can be *three times* as large as the median response.

Table 1: Ratio of Impulse Responses (10th Quantile to Median) Over First Year.

	Monetary Policy Shock	*Credit Spread Shock*
Average Ratio	2.6	2.7
Maximum Ratio	2.9	3.0

This means that a contractionary shock not only makes average (median) outcomes worse, it moves probability mass to the left tail of the GDP growth distribution even relative to what one would expect from median outcomes. This statement holds for both monetary policy and credit spread shocks. These results stress the importance of assessing how shocks move the entire distribution of future outcomes, and not just measures of central tendency.

As we show in the next section, a linear model cannot generate this asymmetry. Hence our

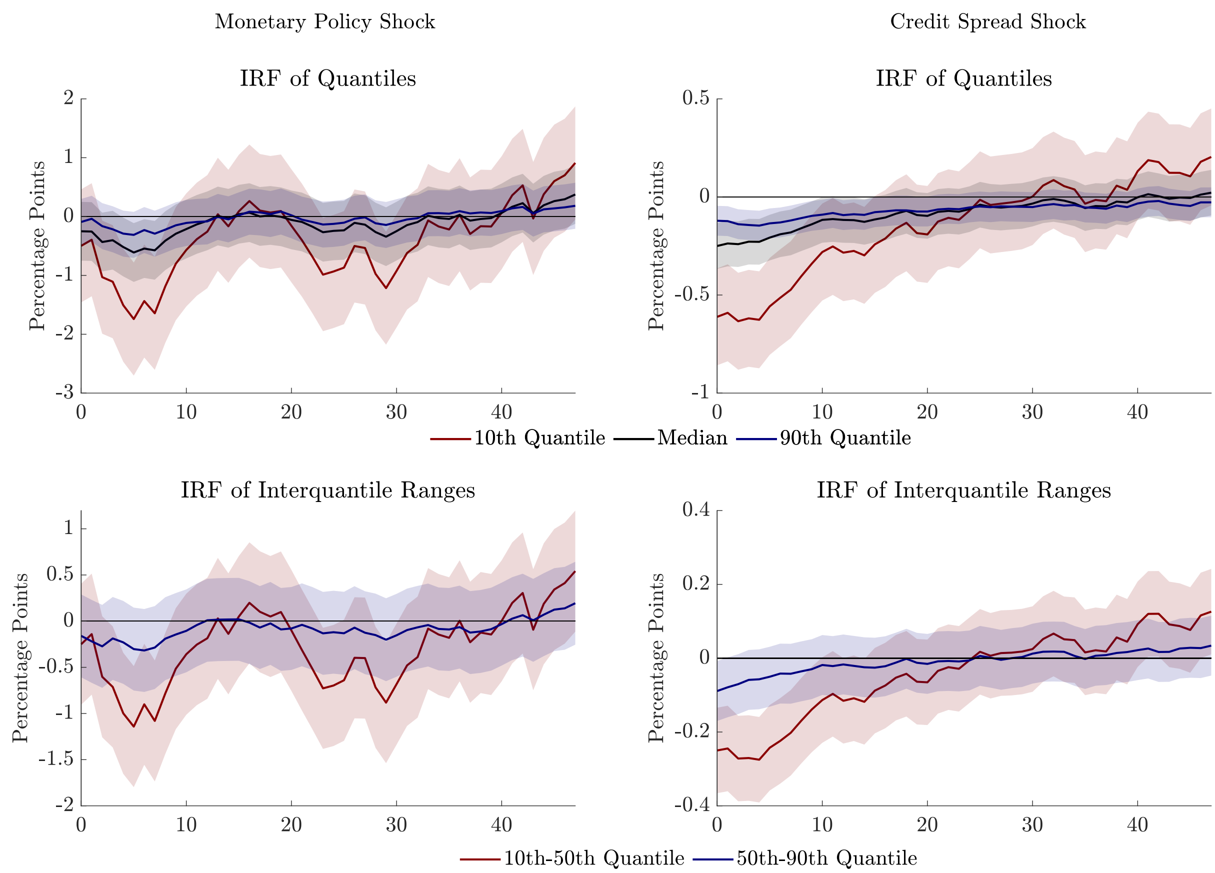

Figure 2: Impulse Responses of Quantiles of Average GDP Growth over the Next Year to *Contractionary* Shocks.

Note: Red is response of the 10th quantile, black is the median response, blue is response of the 90th quantile. Confidence bands correspond to median response, 68% significance level, based on Newey-West and block-of-blocks bootstrapped standard errors. The x-axis gives the horizon of the impulse response, in months. The response on the y-axis is measured in percentage points.

results emphasize that nonlinearities are important to understand macro risk.[13] Furthermore, these nonlinearities cannot be tightly linked to the response to one specific shock, as it is present in both shocks we study. Our results imply that these asymmetries are a feature of the propagation mechanism of shocks, not the original impact. In the next section we describe two data-generating processes that rationalize the asymmetries we find in the data.

[13]As highlighted in Koenker (2005), even linear quantile regressions such as those used here can capture nonlinearities because each quantile is modelled separately.

4 Inspecting the Economic Mechanism

As noted in our introduction, there are at least two key competing theories on the sources of Growth-at-Risk. The first is that some shock has an independent, nonlinear effect on growth. The second is that shocks share a common, nonlinear propagation mechanism.

In the previous section we established that macroeconomic risk is not driven by one shock only. Indeed, we found that both monetary policy and financial shocks generate an asymmetric response in the quantiles of future real GDP growth. In this section, we shed light on potential economic mechanisms underlying our empirical findings. As a by-product, we also validate our two-step approach. To do so, we perform Monte Carlo experiments that draw on a "semi-structural" VAR model with switching coefficients and volatilities, as well as on the nonlinear DSGE model of Gertler, Kiyotaki and Prestipino (2019) featuring bank panics and financial accelerator mechanisms.

4.1 A "Semi-Structural" Threshold VAR Model

We start by considering a specific version of a "semi-structural" threshold VAR model which features the following key mechanism: Whenever financial conditions or macroeconomic conditions worsen considerably, the shock's effect on these variables is more pronounced. This reflects the idea that during recessions nonlinearities arise which exacerbate the effect of adverse shocks.

Model I *Our shock of interest $shock_t$ has a nonlinear effect on the growth outlook through a common propagation mechanism governed by f_t and m_t.*

The VAR model consists of three endogenous variables and three innovations:

$$y_t = \beta_0 + \beta_1 f_t + \beta_2 m_t + \sigma_y e_t^y \tag{4.1}$$

$$f_t = \alpha_1 f_{t-1} + \alpha_2 m_t + \alpha_3(f_{t-1}, m_{t-1}) shock_t + e_t^f \tag{4.2}$$

$$m_t = \gamma_1 m_{t-1} + \gamma_2 f_{t-1} + \gamma_3(f_{t-1}, m_{t-1}) shock_t + e_t^m \tag{4.3}$$

$$\alpha_3(f_t, m_t) = \begin{cases} 5, & \text{if } f_t > f^* \;\&\; m_t < m^* \\ 1, & \text{normal state} \end{cases}, \quad \gamma_3(f_t, m_t) = \begin{cases} -5, & \text{if } f_t > f^* \;\&\; m_t < m^* \\ -1, & \text{normal state} \end{cases}$$

where y_t, f_t, m_t respectively denote GDP growth, a financial factor, and a macroeconomic factor and e_t^y, e_t^f and e_t^m are shocks to these variables; $shock_t$ is the structural shock of interest for the impulse response functions. All shocks are drawn from an independent standard normal distribution. Notice that $\alpha_3(f_t, m_t)$ and $\gamma_3(f_t, m_t)$ are a function of f_t and m_t. In particular, when f_t is larger than a threshold f^* and m_t is smaller than a threshold m^* the effect of an adverse shock on future GDP growth is more pronounced. This captures the idea that due to amplification mechanisms, adverse shocks that hit during already bad times, as measured by tight financial conditions and weak macroeconomic activity, make the growth outlook particularly vulnerable. The other parameter values can be found in Appendix F.
This model encodes the view that some recessions (i.e. situations where m_t is low) can be more severe than others - for instance, if the financial variable f_t is not above the threshold the associated recession tends to be milder. A similar idea is pursued in Jordà, Schularick and Taylor (2020). Even though we focus on one structural shock here for parsimony, our results would be qualitatively the same if instead we had more than one structural shock (i.e. if we replaced $\alpha_3 shock_t$ and $\gamma_3 shock_t$ with $\sum_{i=1}^{I} \alpha_3^i shock_t^i$ and $\sum_{i=1}^{I} \gamma_3^i shock_t^i$, respectively, where $shock_t^i$ is the time t realization of the ith structural shock of interest and the number of structural shocks I is larger than 1).

Impulse Responses We simulate the model 1000 times for 408 periods (the number of periods between January 1986 and December 2019) and store $\{y_t, f_t, m_t, shock_t\}$ for each simulation. We first construct average future GDP growth over the next four periods and estimate its quantiles conditional on the financial and the macroeconomic factor. In Figure 3 we report the impulse responses of the quantiles to the (contractionary) shock of interest, computed via local projection as in our empirical exercise. We again normalize the responses such that, on impact, the median drops by 25 basis points.

There are two important results that emerge. The first is that in the case of the threshold VAR model (left panel), the shock pushes down the left tail more strongly than other parts of the distribution, as in our empirical findings. This can be explained by the fact that when the shock hits during times of financial and macroeconomic distress, its effect is more pronounced, thus making the growth outlook more vulnerable. The second result is that, not surprisingly but reassuringly, no asymmetry is found when we shut down non-linearities in the DGP and thus consider a linear VAR model (right panel). The parameter values for the linear model can also be found in Appendix F.

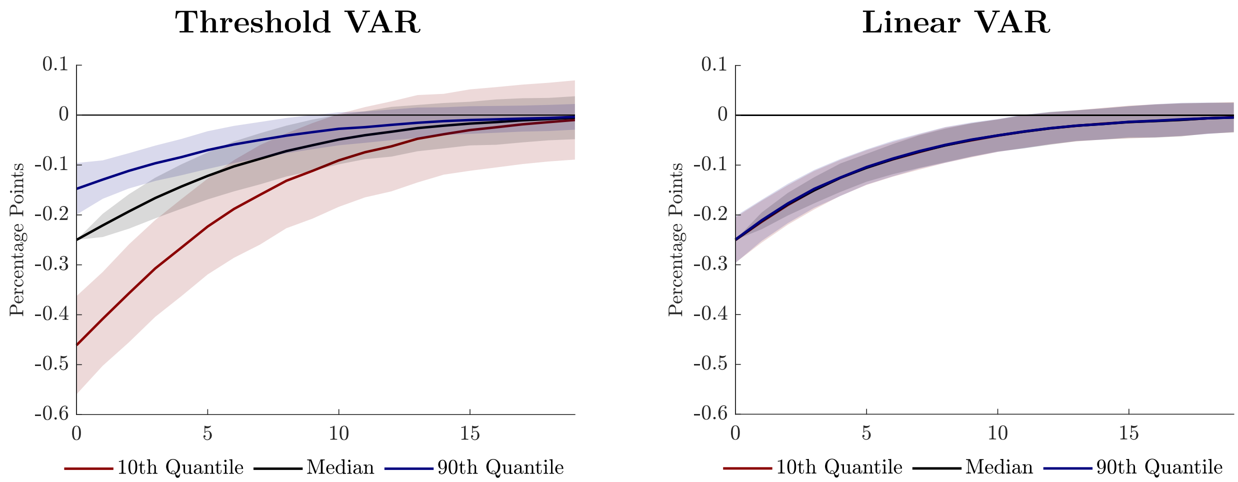

Figure 3: Impulse Responses of Quantiles of Simulated Future GDP Growth from Threshold VAR (Left Panel) and Linear VAR (Right Panel) Model.

Note: Straight lines are medians across simulations. Shaded areas are 68% confidence bands.

In Figure 4 we compare the IRFs of the two-step and one-step approach, in which the responses are obtained by successively leading the dependent variable of the quantile regression (see equation (4.10)). While the responses are similar in shape and magnitude, the one-step approach suffers from higher uncertainty, reason why we favor the two-step procedure in our analysis. Notice that as discussed in Section 2.2, we do not expect the point estimates across approaches to be the same. What is important is that qualitatively the patterns are similar.

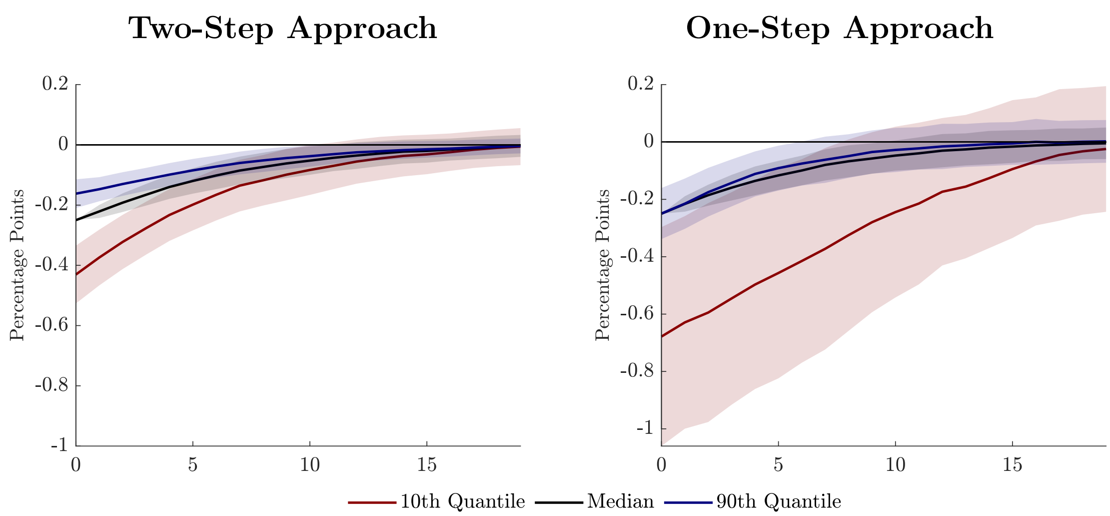

Figure 4: Impulse Responses of Quantiles of Simulated Future GDP Growth from Threshold VAR Model. Two-Step vs. One-Step Approach.

Note: Straight lines are medians across simulations. Shaded areas are 68% confidence bands.

Alternative Specifications of the DGP At this point, it is key to recognize and understand that several data-generating processes can deliver the results shown above. This happens, for instance, in the competing model II below. In this model the shock has a linear effect on financial and macro conditions but the propagation mechanism is such that the latter are nonlinearly related to GDP growth. In particular, it features the following mechanisms. First, whenever financial conditions or macroeconomic conditions worsen, their effect on future GDP growth is more pronounced and so is its variance. (This is governed by the terms $\beta_1(f_t, m_t)$ and $\beta_2(f_t, m_t)$ in the first equation.) This reflects the idea that during recessions nonlinearities arise which can prolong and exacerbate downturns as well as increase uncertainty. Second, both macroeconomic and financial conditions are still driven by a common shock (structural shocks affect multiple aggregates at the same time), but this shock has a linear effect on these variables in this model.

Model II *Shocks affects f_t and m_t linearly, but their effect on y_t is larger during bad times.*

$$y_t = \beta_0 + \beta_1(f_t, m_t) f_t + \beta_2(f_t, m_t) m_t + \sigma_y(f_t, m_t) e_t^y \tag{4.4}$$

$$f_t = \alpha_1 f_{t-1} + \alpha_2 m_t + \alpha_3 shock_t + e_t^f \tag{4.5}$$

$$m_t = \gamma_1 m_{t-1} + \gamma_2 f_{t-1} + \gamma_3 shock_t + e_t^m \tag{4.6}$$

$$\beta_1(f_t, m_t) = \begin{cases} -1.5, & \text{if } f_t > f^* \ \& \ m_t < m^* \\ -0.5, & \text{normal state} \end{cases}, \quad \beta_2(f_t, m_t) = \begin{cases} 1.5, & \text{if } f_t > f^* \ \& \ m_t < m^* \\ 0.5, & \text{normal state} \end{cases},$$

$$\sigma_y(f_t, m_t) = \begin{cases} 2, & \text{if } f_t > f^* \ \& \ m_t < m^* \\ 0.1, & \text{normal state} \end{cases}$$

If we again simulate data from this competing model, feed it to our two-step procedure and compute the impulse responses, the model would again signal that the effect of the shock on GDP growth is nonlinear – as evident in the right panel of Figure 5. In particular, we obtain again the result of model I, in the left panel, that the lower tail moves down more than the median. This result is not surprising since financial and macro conditions are the key variables governing the nonlinearity of future growth in our threshold VAR and both of these variables are in turn driven by that common shock.

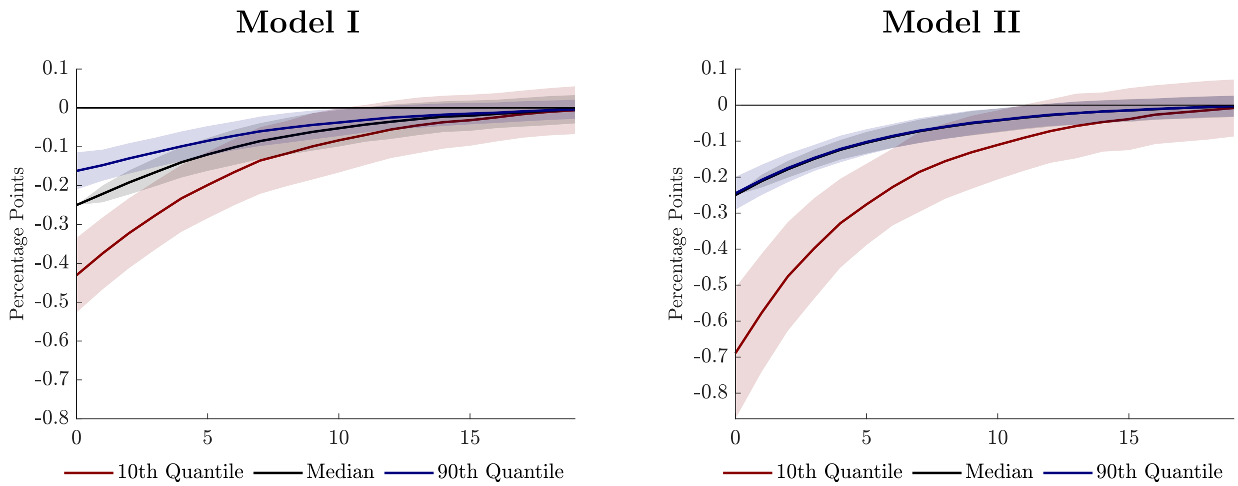

Figure 5: Comparison of Impulse Responses of Quantiles of Simulated Future GDP Growth from Threshold VAR. Model I vs. Model II.

Note: Straight lines are medians across simulations. Shaded areas are 68% confidence bands.

These same patterns can be replicated by other DGPs, like model III below, in which the asymmetric effect of the shock is propagated nonlinearly through macroeconomic activity m_t only. However, because the asymmetry with respect to m_t is inherited by the financial side (since f_t depends on m_t), only the impulse response of f_t to $shock_t$ would unveil the true source of the nonlinearity (m_t). Model III thus acknowledges how the interdependence between macro and financial conditions can confound inference on which variables are responsible for the asymmetric response of GDP growth to shocks.

Model III *Asymmetry from macroeconomic conditions m_t is inherited by f_t and y_t.*

$$y_t = \beta_0 + \beta_1 f_t + \beta_2 m_t + \sigma_y e_t^y \tag{4.7}$$

$$f_t = \alpha_1 f_{t-1} + \alpha_2 m_t + \alpha_3 shock_t + e_t^f \tag{4.8}$$

$$m_t = \gamma_1 m_{t-1} + \gamma_2 f_{t-1} + \gamma_3(m_{t-1}) shock_t + e_t^m \tag{4.9}$$

Discriminating Theories Some DGPs can be discriminated by looking at the response of the variables determining GDP growth (f_t and m_t) to the shock. In model I, the nonlinear relationship between shock and growth, is a result of the shock's asymmetric effect on financial and macro conditions themselves. This is confirmed by their impulse responses in the top panel of Figure 6, which we compute in one step. Contractionary shocks lead to more vulnerable financial conditions, with the 90th quantile moving up more than the median, as well macroeconomic activity (with the 10th quantile dropping more than the median).

This result cannot be established in model II, where the shock has a symmetric effect on financial and macro conditions, as illustrated in the bottom panel of Figure 6. Likewise, in model III only macro conditions respond asymmetrically to the shock, since the financial factor is linearly related to it. Looking at the response of the unconditional distribution of the conditioning variables of our quantile regression model can provide valuable insights into the mechanism driving the asymmetric response of the growth outlook to structural shocks.

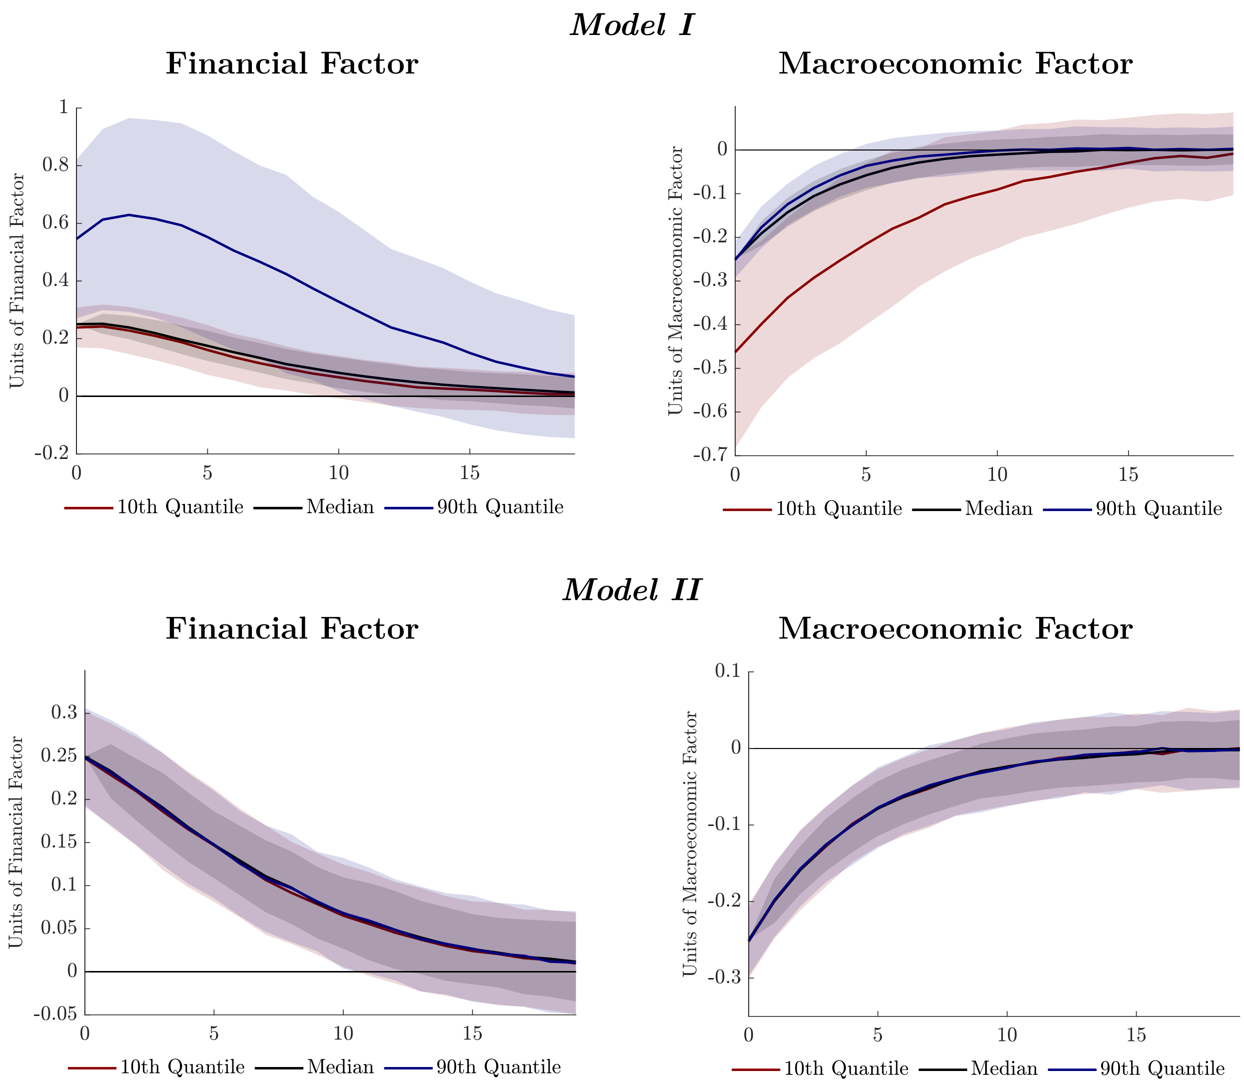

Figure 6: Impulse Responses of Simulated Financial Factor and Macroeconomic Factor from Threshold VAR. Model I vs. Model II.

Note: Straight lines are medians across simulations. Shaded areas are 68% confidence bands.

We thus now focus on the impulse responses of the financial and macro factor quantiles and interquantile ranges using US data. These responses are estimated in one step, where the quantiles are estimated via:

$$q_{\tau}\left(X_{t+s}\right) = \delta_{\tau}^{s} + \theta_{\tau}^{s}\ shock_t + \Psi(L)_{\tau}^{s}\ controls_t \quad s = \{0, \ldots, S\} \tag{4.10}$$

where $X_t = \{FF_t,\ MF_t\}$ and $q_{\tau}\left(X_{t+s}\right)$ is the quantile τ of X_t at horizon s. The sequence of coefficients θ_{τ}^{s} traces out the impulse response to a structural shock.

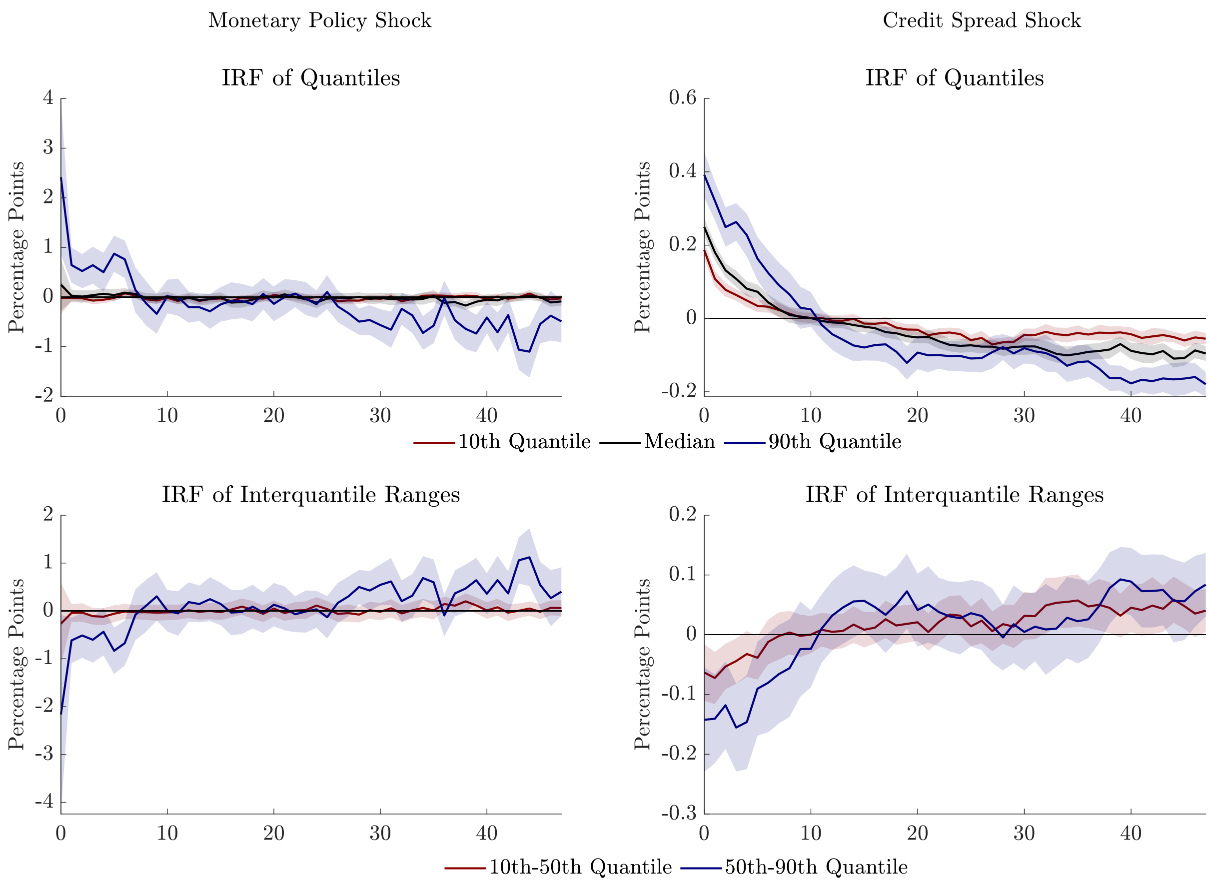

Figure 7: Impulse Responses of Quantiles of Financial Factor to *Contractionary* Shocks using US Data.

Note: Red is response of the 10th quantile, black is the median response, blue is response of the 90th quantile. Confidence bands correspond to median response, 68% significance level, based on block-of-blocks bootstrapped standard errors. For the interquantile ranges, the bands are obtained by summing the uncertainty of the quantile-pair considered (since we do not have an estimate for the covariance) and are thus on the conservative end. The x-axis gives the horizon of the impulse response, in months. The response on the y-axis is measured in percentage points.

These IRFs are crucial in giving insight into the nature of the nonlinear relationship between the shock and the growth outlook. The results in Figures 7 and 8, respectively for the financial and macro factor, point to an asymmetric response of both variables to adverse shocks. In particular, the monetary policy and the financial shocks make the growth outlook more vulnerable in that they disproportionately push up the upper tail of the financial factor (in blue) and bring down the lower tail of the macroeconomic factor (in red).

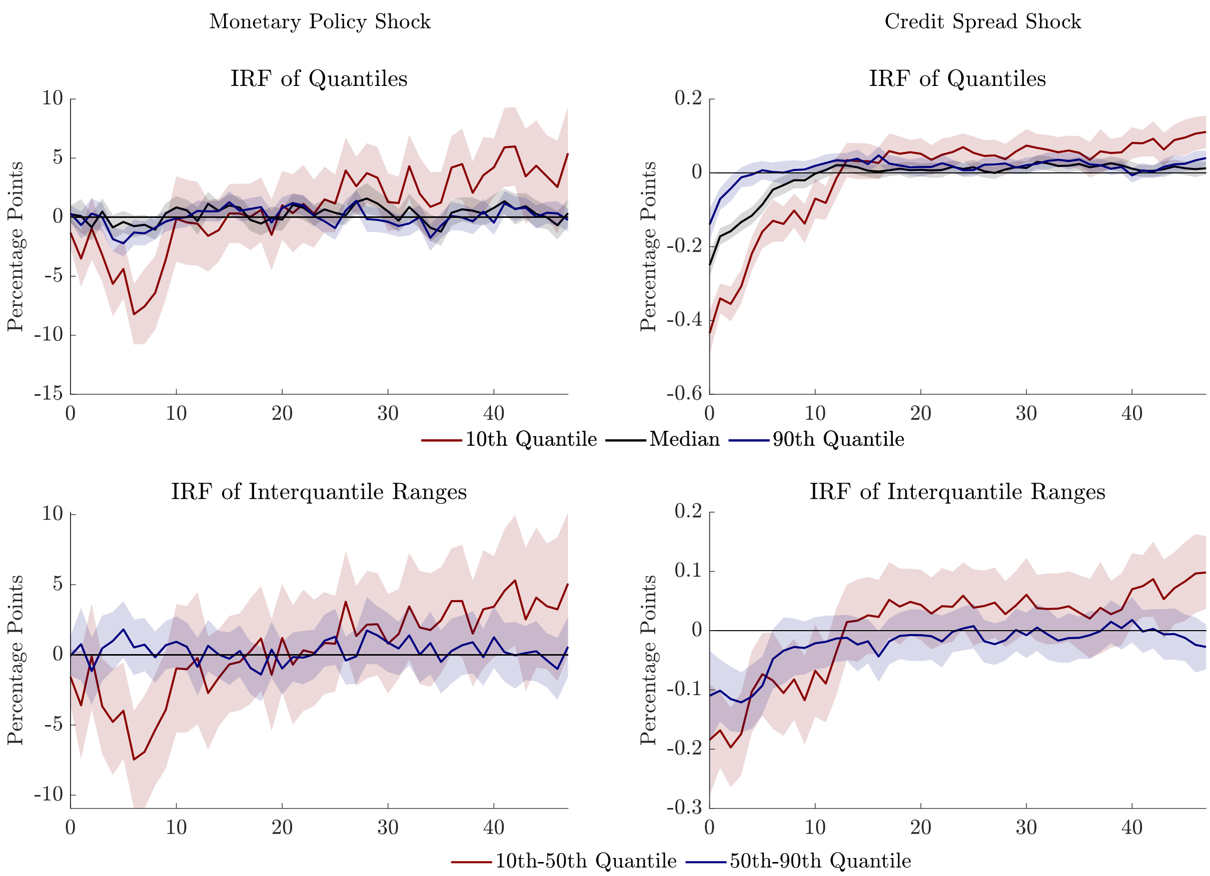

Figure 8: Impulse Responses of Quantiles of Macroeconomic Factor to *Contractionary* Shocks using US Data.

Note: Red is response of the 10th quantile, black is the median response, blue is response of the 90th quantile. Confidence bands correspond to median response, 68% significance level, based on block-of-blocks bootstrapped standard errors. For the interquantile ranges, the bands are obtained by summing the uncertainty of the quantile-pair considered (since we do not have an estimate for the covariance) and are thus on the conservative end. The x-axis gives the horizon of the impulse response, in months. The response on the y-axis is measured in percentage points.

4.2 A Macroeconomic Model with Financial Panics

Next, we show that the nonlinear DSGE model of Gertler, Kiyotaki and Prestipino (2019) features the same properties as Model I and thus can also rationalize our empirical findings.

The Model We briefly discuss the key mechanism that generates the nonlinearity. We refer interested readers to the original paper for details of the model.[14] The model is fully microfounded and extends the conventional New Keynesian model with investment by introducing bankers. Bankers are more efficient than households in handling loans. However, bankers are constrained in their ability to raise external funds and are subject to runs. The latter gives rise to multiple equilibria: one with and one without a financial panic.

A financial panic forces the banking system into liquidation, expanding the share of capital held by households. The reallocation of capital holdings from bankers to less efficient households increases the cost of capital, which ultimately disrupts firms' borrowing. Consequently, investment drops substantially more than in the equilibrium without a bank run. A self-fulfilling financial panic equilibrium exists if and only if, in the event of all other depositors' run, an individual household will be better off to follow the run. When financial conditions are strong, the economy fluctuates around a standard equilibrium. In contrast, when the financial system is weak (i.e., at the edge of the bank-run regime), even a small shock can push the economy into a self-fulfilling bank-run equilibrium. Combined with a sunspot shock, this triggers a financial panic and a deep recession.

Model Solution The model is highly non-linear, and the non-linear effects of structural shocks to the real economy depend on financial conditions. To allow for these non-linear transition dynamics, the model is solved non-linearly using the collocation method with policy functions solved by time iteration. We follow the original paper in focusing on a capital quality shock as a representative structural shock. However, other shocks would give rise to qualitatively very similar results because the mechanism creating asymmetry is not specifically tied to one structural shock.

Parameterization We simulate the model using the original calibration of the deep parameters and of the capital quality shock process (the only fundamental shock in the model). In order to generate rare financial crisis, we calibrate the process for the sunspot shock such that

[14]The full model consists of 40 equations which we do not include here for the sake of brevity.

a bank run equilibrium arises after a big negative shock (above two standard deviations).[15]

Simulation We simulate this model 1000 times for 408 periods (the number of periods between January 1986 and December 2019) and store the level of GDP, the credit spread, and the capital quality shock. In our analysis, we treat the credit spread as the equivalent of the financial factor though results are robust to the use of alternative measures of financial conditions in the model. The simulated data shown in the top panel of Figure 9 indicate that also in this model there is a non-linear relationship between growth and financial conditions. Indeed, large credit spreads are associated with extremely negative growth realizations.

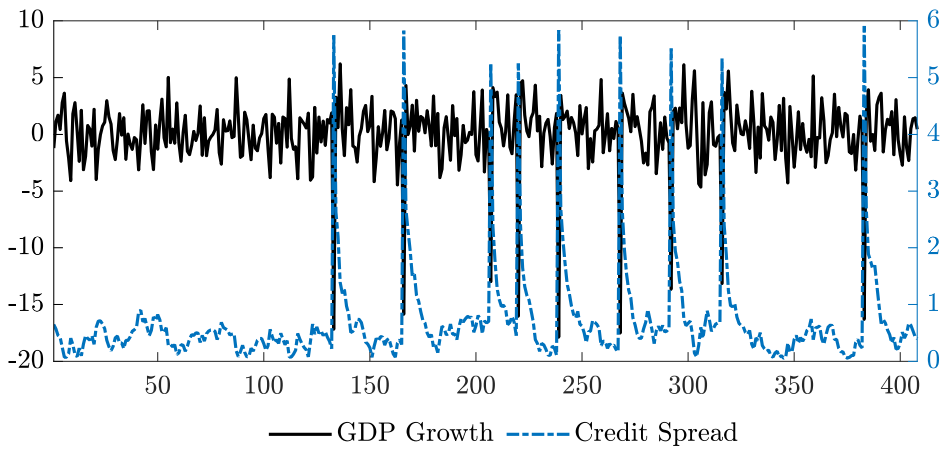

Note: Example from one simulation. GDP Growth (left axis) and Credit Spread (right axis).

Figure 9: Simulated Data from Gertler, Kiyotaki and Prestipino (2019) Model.

Quantile Regression We focus on current (period-over-period) GDP growth as the shocks are transitory and only create a sharp but short-lived recession.[16] Further, we only consider the credit spread as conditioning variable since the non-linearity in the model is mainly coming from financial conditions. As shown in Table 2, the quantile regression picks up the

[15] Our calibration intends to make the bank run event (see Figure 2 of the original paper) re-occurring in the simulated sample. In their event study, the authors feed in two consecutive negative capital quality shocks of roughly one standard deviation to push the economy to the edge of the bank-run regime. A third shock, together with a sun-spot shock, then produce the bank run event. We deviate slightly from their event study by having sun-spot shocks occurring concurrently with a negative two standard deviations capital quality shocks to ensure the probability of a bank run of 2.5% across simulated samples.

[16] Notice that even though the distribution is for current GDP growth, this exercise is still relevant for the characterization of risks to the growth outlook from a forecasting perspective, as data on credit spreads is available at a higher frequency than GDP growth data in the US.

nonlinear relationship between financial conditions and economic growth, suggesting that financial conditions have a more negative effect on the left tail than on the median.

Table 2: Quantile Regression Coefficients Estimated on Simulated Data from Gertler, Kiyotaki and Prestipino (2019) Model.

Variable	10th Quantile	Median	90th Quantile
Intercept	-0.94	0.91	2.55
Credit Spread	-2.64	-1.32	0.45

* Note: Average Across Simulations.

Thus, also in this example, the left tail of the distribution falls substantially during times of extreme financial distress, characterizing a vulnerable growth outlook. This becomes evident in Figure 10 which plots the resulting quantiles of GDP growth.

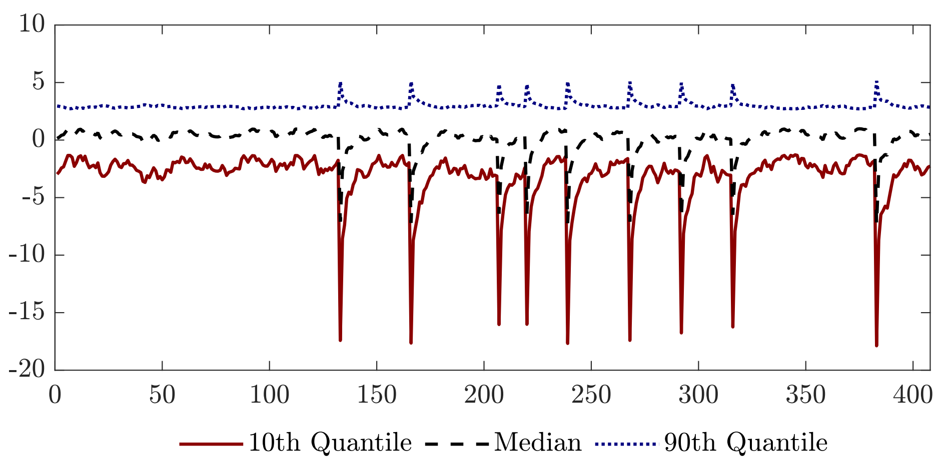

Note: Example from one simulation.

Figure 10: GDP Growth Quantiles from Gertler, Kiyotaki and Prestipino (2019).

Impulse Responses Figure 11 presents the impulse responses of the quantiles to the (contractionary) capital quality shock, computed as in our empirical exercise. Also in this example, we normalize the responses such that, on impact, the median drops by 25 basis points. We again find that the shock pushes down the left tail more strongly than other parts of the distribution, as in our empirical findings. Indeed, while the shock enters the model linearly, as in the previous exercise, its relationship to GDP growth is nonlinear as it affects the financial variables that transmit and amplify the disturbance nonlinearly.

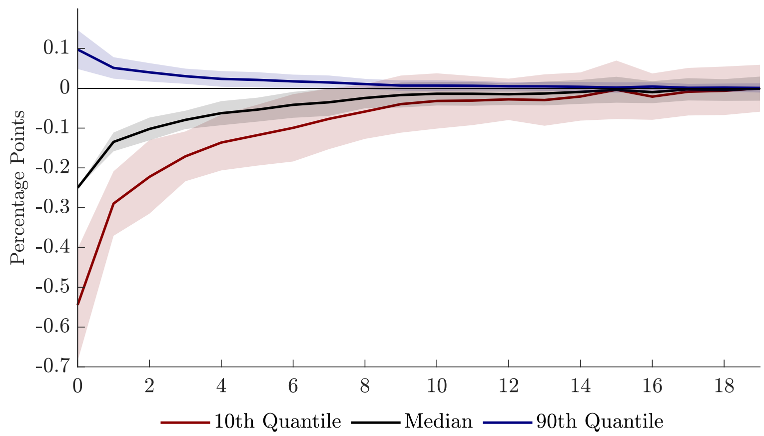

Figure 11: Impulse Responses of Quantiles of Simulated GDP Growth to a Capital Quality Shock from Gertler, Kiyotaki and Prestipino (2019) Model.

Note: Straight lines are medians across simulations. Shaded areas are 68% confidence bands.

Crucially, we now ask how the key variable through which shocks have a nonlinear effect on GDP growth in the model, namely credit spreads, responds to the capital quality shock. In Figure 12, an adverse capital quality shock pushes up the upper quantile of credit spreads disproportionately more. This replicates a key finding from our empirical exercise and a key pattern of the threshold VAR model which we considered previously: That is, adverse shocks put GDP growth more at risk by making financial conditions more vulnerable.

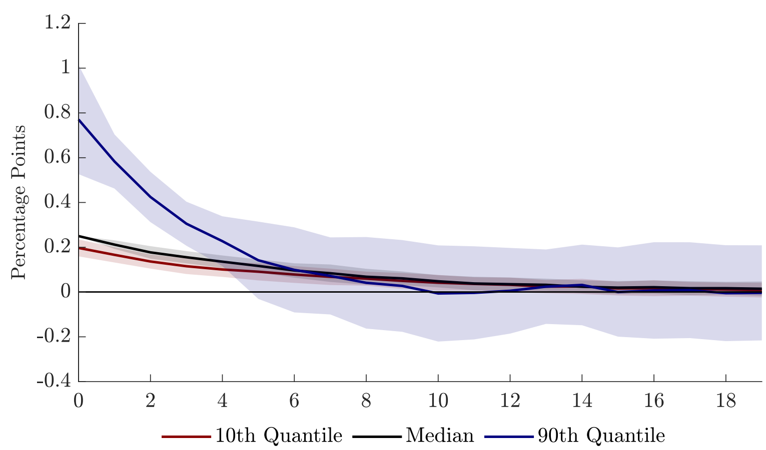

Figure 12: Impulse Responses of Quantiles of Simulated Credit Spread to a Capital Quality Shock from Gertler, Kiyotaki and Prestipino (2019) Model.

Note: Straight lines are medians across simulations. Shaded areas are 68% confidence bands.

5 Conclusion

This paper studies tail risk in U.S. aggregate outcomes. In particular, we study how macroeconomic shocks affect tail risk. All shocks considered (monetary policy and financial shocks) affect tail risk disproportionately more than other quantiles (on average the response of the 10th percentile over the first year after impact is around 3 times larger than the median). Contractionary shocks thus deserve even more attention than what their effect on average outcomes suggests to the extent that they make poor economic conditions more likely. Our results suggest that policymakers should be especially weary of unexpected adverse changes in the economy.

The fact that all shocks we study display this tail risk asymmetry points to a common mechanism lying behind these asymmetries. Indeed, the two data-generating processes we use as laboratories to test this hypothesis can replicate our findings: In the first experiment, a threshold VAR model, a common mechanism (which could be thought of as a financial accelerator mechanism in a more structural model) amplifies the negative effect on economic growth of a deterioration in financial and macroeconomic conditions. Guided by insights from this model, we show also empirically that adverse shocks disproportionately increase financial stress and macroeconomic strain in bad times. In the second, a macroeconomic model with financial panics, the bank-run equilibrium features nonlinearities that magnify the effects of a financial panic on economic activity. In that sense, we find that nonlinear models are promising for modeling and studying macroeconomic risk.

In providing a conceptual framework to think about the sources of macroeconomic risk, our paper provides insights on how to replicate growth-at-risk patterns in theoretical models with a particular focus on impulse responses to structural shocks.

References

Adrian, T., Boyarchenko, N. and Giannone, D. (2019). Vulnerable Growth. *American Economic Review*, **109** (4), 1263–1289.

—, Liang, N., Zabczyk, P. and Duarte, F. (2020). *Monetary and Macroprudential Policy with Endogenous Risk*. IMF Working Papers 2020/236, International Monetary Fund.

Aikman, D., Bluwstein, K. and Karmakar, S. (2021). *A tail of three occasionally-binding constraints: a modelling approach to GDP-at-Risk*. Bank of England working papers 931, Bank of England.

Antolín-Díaz, J. and Rubio-Ramírez, J. F. (2018). Narrative Sign Restrictions for SVARs. *American Economic Review*, **108** (10), 2802–29.

Barnichon, R. and Brownlees, C. (2019). Impulse Response Estimation by Smooth Local Projections. *The Review of Economics and Statistics*, **101** (3), 522–530.

Caldara, D., Cascaldi-Garcia, D., Cuba Borda, P. and Loria, F. (2020). Understanding Growth-at-Risk: A Markov-Switching Approach. *mimeo*.

Carriero, A., Clark, T. E. and Marcellino, M. (2020). *Capturing Macroeconomic Tail Risks with Bayesian Vector Autoregressions*. Working Papers 202002, Federal Reserve Bank of Cleveland.

Chavleishvili, S. and Manganelli, S. (2017). Quantile Impulse Response Functions.

Cogley, T. and Sargent, T. J. (2005). Drift and Volatilities: Monetary Policies and Outcomes in the Post WWII U.S. *Review of Economic Dynamics*, **8** (2), 262–302.

Cúrdia, V., Del Negro, M. and Greenwald, D. L. (2014). Rare Shocks, Great Recessions. *Journal of Applied Econometrics*, **29** (7), 1031–1052.

Fagiolo, G., Napoletano, M. and Roventini, A. (2008). Are Output Growth-Rate Distributions Fat-Tailed? Some Evidence from OECD Countries. *Journal of Applied Econometrics*, **23** (5), 639–669.

Favara, G., Gilchrist, S., Lewis, K. F. and Zakrajšek, E. (2016). Updating the Recession Risk and the Excess Bond Premium. (2016-10-06).

Gertler, M. and Karadi, P. (2015). Monetary Policy Surprises, Credit Costs, and Economic Activity. *American Economic Journal: Macroeconomics*, **7** (1), 44–76.

—, Kiyotaki, N. and Prestipino, A. (2019). A Macroeconomic Model with Financial Panics. *The Review of Economic Studies*, **87** (1), 240–288.

Giglio, S., Kelly, B. and Pruitt, S. (2016). Systemic risk and the macroeconomy: An empirical evaluation. *Journal of Financial Economics*, **119** (3), 457 – 471.

Gilchrist, S. and Zakrajšek, E. (2012). Credit Spreads and Business Cycle Fluctuations. *American Economic Review*, **102** (4), 1692–1720.

Jarociński, M. and Karadi, P. (2020). Deconstructing Monetary Policy Surprises-the role of Information Shocks. *American Economic Journal: Macroeconomics*, **12** (2), 1–43.

Jordà, O. (2005). Estimation and Inference of Impulse Responses by Local Projections. *American Economic Review*, **95** (1), 161–182.

Jordà, O., Schularick, M. and Taylor, A. M. (2020). *Disasters Everywhere: The Costs of Business Cycles Reconsidered*. Working Paper 26962, National Bureau of Economic Research.

Justiniano, A. and Primiceri, G. E. (2008). The Time-Varying Volatility of Macroeconomic Fluctuations. *American Economic Review*, **98** (3), 604–641.

Kilian, L. and Manganelli, S. (2008). The Central Banker as a Risk Manager: Estimating the Federal Reserve's Preferences under Greenspan. *Journal of Money, Credit and Banking*, **40** (6), 1103–1129.

Kim, T.-H., Lee, D. J. and Mizen, P. (2019). Impulse Response Analysis in Conditional Quantile Models and an Application to Monetary Policy.

Koenker, R. (2005). *Quantile Regression*. Econometric Society Monographs, Cambridge University Press.

— and Bassett, G. (1978). Regression Quantiles. *Econometrica*, **46** (1), 33–50.

Linnemann, L. and Winkler, R. (2016). Estimating Nonlinear Effects of Fiscal Policy Using Quantile Regression Methods. *Oxford Economic Papers*, **68** (4), 1120–1145.

Mertens, K. and Ravn, M. O. (2013). The Dynamic Effects of Personal and Corporate Income Tax Changes in the United States. *American Economic Review*, **103** (4), 1212–1247.

Miranda-Agrippino, S. and Ricco, G. (2020). The Transmission of Monetary Policy Shocks. *American Economic Journal: Macroeconomics*, **forthcoming**.

Mumtaz, H. and Surico, P. (2015). The Transmission Mechanism In Good And Bad Times. *International Economic Review*, **56**, 1237–1260.

Nakamura, E. and Steinsson, J. (2018). High-frequency Identification of Monetary Non-neutrality: the Information Effect. *The Quarterly Journal of Economics*, **133** (3), 1283–1330.

Newey, W. K. and West, K. D. (1987). A Simple, Positive Semi-Definite, Heteroskedasticity and Autocorrelation Consistent Covariance Matrix. *Econometrica*, **55** (3), 703–708.

Plagborg-Møller, M., Reichlin, L., Ricco, G. and Hasenzagl, T. (2020). *When is Growth at Risk?* Tech. rep., Brooking Papers on Economic Activity.

— and Wolf, C. (2019). Local Projections and VARs Estimate the Same Impulse Responses.

Primiceri, G. E. (2005). Time Varying Structural Vector Autoregressions and Monetary Policy. *Review of Economic Studies*, **72** (3), 821–852.

Ramey, V. (2016). Macroeconomic Shocks and Their Propagation. *Handbook of Macroeconomics*, **2**, 71 – 162.

Romer, C. D. and Romer, D. H. (2004). A New Measure of Monetary Shocks: Derivation and Implications. *American Economic Review*, **94** (4), 1055–1084.

Uhlig, H. (2005). What Are the Effects of Monetary Policy on Output? Results from an Agnostic Identification Procedure. *Journal of Monetary Economics*, **52** (2), 381–419.

Zhang, D. (2019). Term Structure, Forecast Revision and the Signaling Channel of Monetary Policy. *mimeo*.

A Some Intuition for Impulse Responses of Quantiles

This section gives three examples where an initial distribution of a variable changes after a shock hits. We show these examples to illustrate how a change in quantiles is linked to changes in the distribution as a whole and how changes in specific moments translate into changes in quantiles.[17] Our scenario is as follows: After an initial distribution of a scalar variable is hit by a shock, we trace out how this distribution changes on impact and in the period after impact. We consider three experiments:

1. The shock leads to an increase in the variance of our distribution, which is Gaussian.
2. The shock leads to an increase in the mean of our distribution, which is Gaussian.
3. The shock leads to a change in *both* the mean and variance of our distribution, which is Gaussian.

Figure A-1 plots three panels for each experiment. The first panel in each row shows the initial distribution, the distribution when the shock hits, and the distribution in the period after the shock has materialized. The middle panel in each row shows the evolution of the 10th and 90th percentile for those three periods. The last panel in each row gives the impulse responses for the 10th and 90th percentiles *under the assumption that if the shock that moved the distributions had not materialized, the distribution would have remained at its original position.*[18] As the impulse response plots the difference between the relevant percentiles and the original values, the impulse response figures only show values for two time periods (the period when the shock hits and the period after). Each row presents the figures for one experiment. Note that the levels of the percentiles are not directly interpretable as IRFs because we do not subtract the baseline value from the quantiles in those figures. As we can see, an increase in the variance of a symmetric distribution makes the quantiles drift apart in a mirror-image fashion, whereas a change in the mean of a symmetric distribution makes the quantiles move in parallel, which in turn makes the impulse responses lie on top of each other. With a non-symmetric distribution (or if a shock makes a distribution non-symmetric) the quantiles can drift apart, but not necessarily in a mirror-image fashion, as highlighted in the third example.

[17] These examples are not meant to be exhaustive. There can be many other changes in distributions that lead to similar movements in quantiles to those displayed in this section.

[18] For our purposes, impulse responses are defined as the difference between two conditional expectations, where one of the expectations conditions on a specific value for one shock in one specific period.

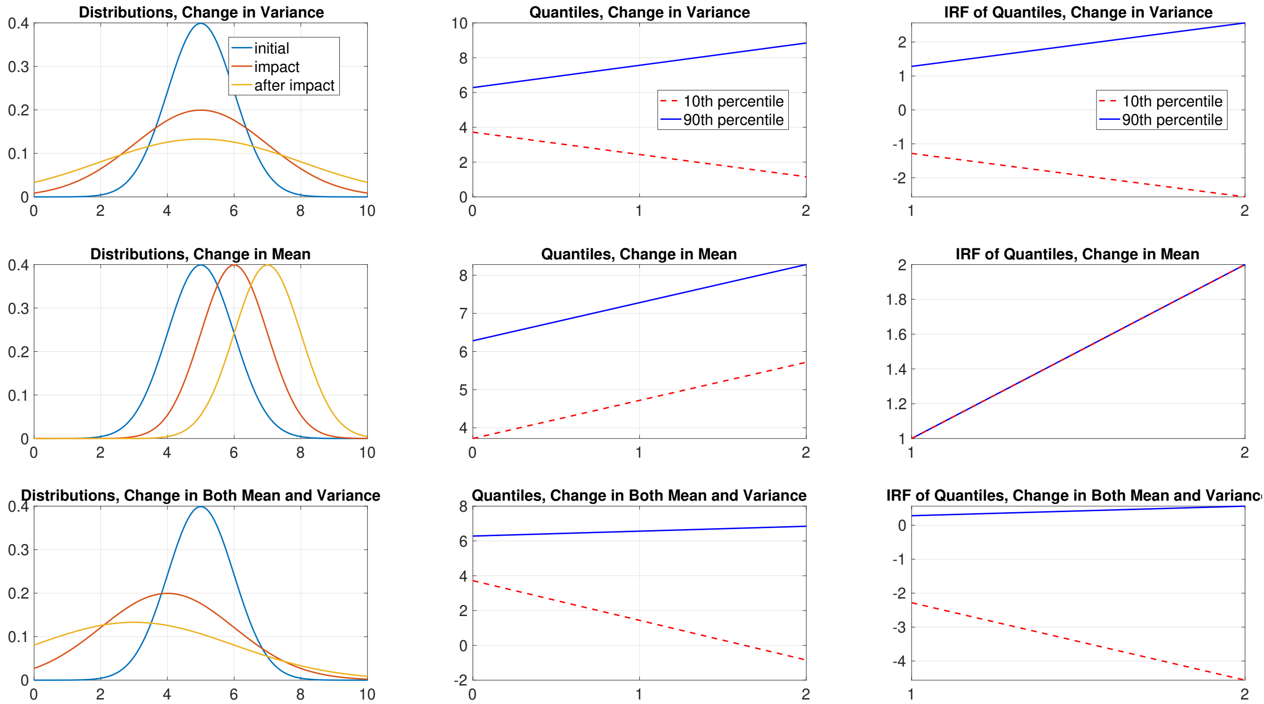

Figure A-1: Illustration of Changes in Percentiles.

Interpreting changes in multiple quantiles jointly can be challenging because we have to envision how the entire distribution changes. As an example, let us focus on the third experiment. As can be seen from the last panel on the bottom row of Figure A-1, the 10th and 90th percentile drift apart because the 90th percentile increases slightly, whereas the 10th percentile decreases substantially. Thus the distribution *spreads out* as a result of the shock—this can also be seen by looking at the leftmost panel of the bottom row, where the yellow distribution is more spread out than the original blue distribution. Let us for a second imagine that this impulse response is the response to a contractionary shock and that quantiles react linearly to those shocks (as will be the case in our local projections). Such shocks would not only move mean and median of the distribution (this can be seen from looking directly at the Gaussian distributions in the bottom left panel, where the mean/median of the Gaussian distribution moves from 5 to 3), but it actually moves the 10th percentile substantially more, thus not only making average outcomes worse, but making outcomes in the left tail much more likely. This is the relevant case in our emprical results.

Another scenario that could occur is that the impulse responses of various quantiles cross. It is important to emphasize that this does *not* mean that the quantiles themselves cross. In fact, a crossing of impulse responses of various quantiles can just mean that one quantile reverts back to its pre-shock value faster than another.

Our responses capture the difference between the expected path of the τth quantile at horizon h after a given shock of specific (one standard deviation) size occurs and the expected path of the τth quantile conditioning on such a specific shock value. That is, a response equal to 0 at horizon h means that the expected τth quantile at horizon h is the same independently of whether we condition on a specific shock value in the initial period.

Next, we show an example where the impulse responses of the quantiles cross. By construction, however, the quantiles themselves cannot cross because we directly model changes in the entire distribution. The example is similar to the first example from the main text, but we change the sequence of variances for the Gaussian distribution to achieve the crossing of the impulse responses of the 10th and 90th percentiles.

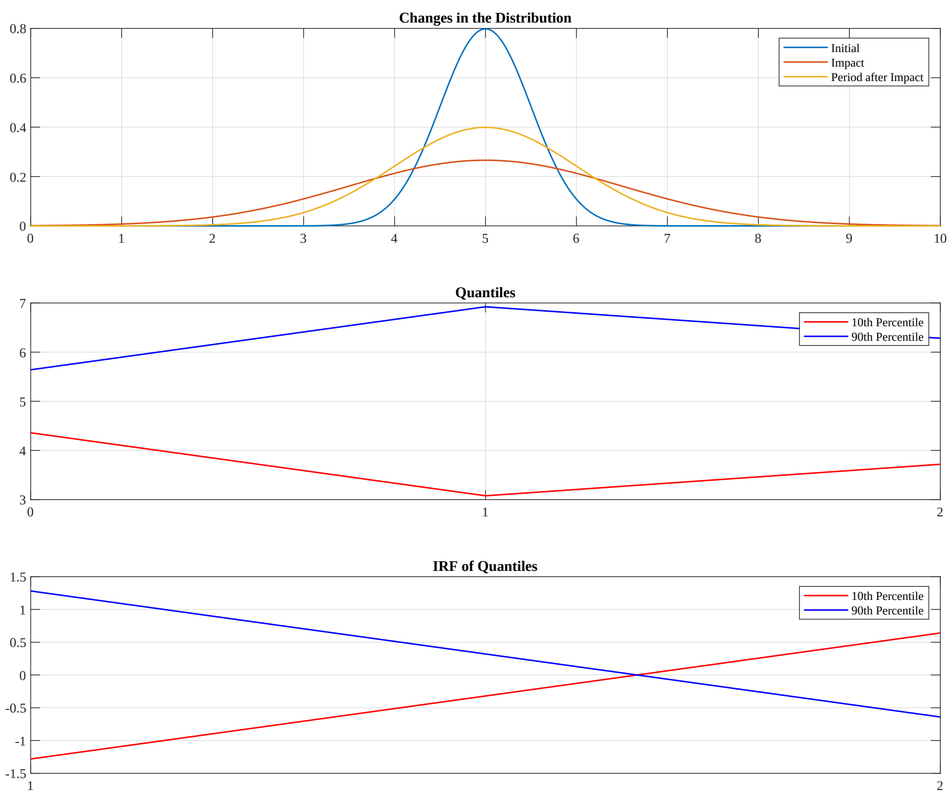

Figure A-2: Example Where Impulse Responses Cross But Quantiles Don't.

B Data

This section gives a brief overview of the data we use throughout this paper.

B.1 Growth-at-Risk Data

The data for the quantile regression com from Caldara, Cascaldi-Garcia, Cuba Borda and Loria (2020). The dependent variable is monthly GDP growth, shown in Figure B-1 whereas the conditioning variables are a financial and a macroeconomic factor, displayed in Figure B-2. These three variables are estimated from a Dynamic Factor Model (DFM). The reader is referred to the original paper for the details on the DFM model and its estimation procedure.

The following set of data inform the two factors used as conditioning variables in the quantile regression:

1. Financial factor
    - Volatility index of the S&P 100 (VXO)
    - Excess bond premium
    - 3-month LIBOR rate minus 3-month Treasury bill
    - 3-month financial commercial paper rate minus 3-month Treasury bill.
2. Macroeconomic factor
    - Industrial Production
    - Retail Sales
    - New Export Orders Component of Purchasing Managers' Index (PMI)
    - Initial Unemployment Claims
    - GDP

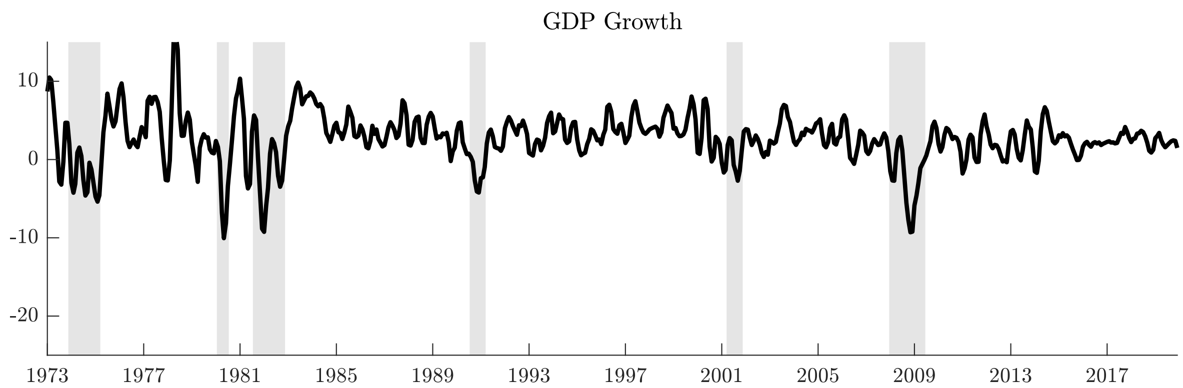

Figure B-1: Monthly Real GDP Growth.

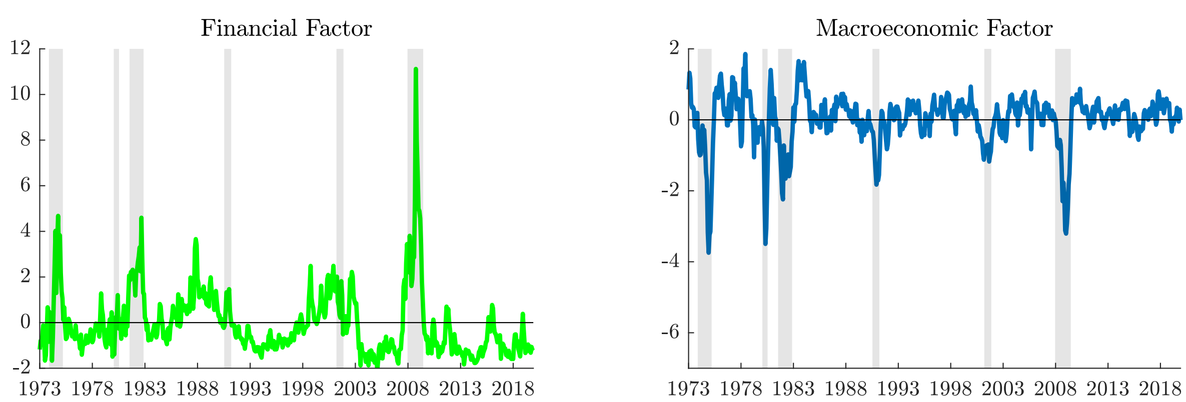

Figure B-2: Conditioning Variables in the Quantile Regressions.

B.2 Local Projection Variables

Most control variables used in the local projection stage are available in FRED. The source data for the control variables are:

- Consumer Price Index for all Urban Consumers: All Items.
  FRED Mnemonic: CPIAUCSL.
- Federal funds rate. FRED Mnemonic: FEDFUNDS.
- Excess bond premium. Source: Favara, Gilchrist, Lewis and Zakrajšek (2016)[19].
- Real GDP. Percent Change from Preceding Period, Seasonally Adjusted Annual Rate.
  FRED Mnemonic: A191RL1Q225SBEA.
- Chicago Fed National Financial Conditions Index. FRED Mnemonic: NFCI.
- Industrial production. FRED Mnemonic: INDPRO.
- One-year government bond rate. FRED Mnemonic: GS1.

As to the shocks used in the robustness exercises, the Romer and Romer (2004) monetary shock is provided by Ramey (2016). We aggregate the monthly shock series to quarterly frequency by taking the quarterly average. We take the narrative monetary policy shock provided by Antolín-Díaz and Rubio-Ramírez (2018), again aggregated to quarterly frequency by calculating the quarterly average.

[19]The series can be downloaded at https://www.federalreserve.gov/econresdata/notes/feds-notes/2016/updating-the-recession-risk-and-the-excess-bond-premium-20161006.html.

C Bootstrap Procedure for Two-Step Approach

This section gives a brief overview of the bootstrap procedure used to obtain the confidence bands. The procedure is designed to capture the uncertainty involved both with the quantile regression and with the local projection step of our estimation approach.

Quantile Regression Step The first step of the bootstrap procedure involves quantifying the uncertainty around the quantile regression estimates. To do so, we use a "blocks-of-blocks" bootstrap.

For a total number of $K = 100$ bootstrap replications, blocks of data are randomly drawn to form a new sample of the same size as the original data. Importantly, the blocks are resampled in the same order for both the dependent variable y and the regressors X, a key step which preserves the time-dependency in the data.

We pass on to the local projection stage, the time series of estimated quantiles associated with each of these resampled data sets. The local projection coefficient $\hat{\theta}_\tau^s$ then captures the effect of the shock on the quantile τ at horizon s, for each bootstrap replication k. The procedure is asymptotically valid for stationary processes if the block size l increases at a suitable rate as $T \to \infty$. We set $m = \sqrt[3]{T}$, where T is the sample size. Finally, this bootstrap procedure preserves the quantile regression feature of being agnostic about the underlying distribution of the error terms, as this is not a residual-based approach.

Local Projection Step At each horizon s, 100 bootstrap replications of the local projection estimates are obtained by drawing the impulse response coefficients from their asymptotic distribution. This distribution is known and given by $\theta_\tau^s \sim \mathcal{N}\left(\hat{\theta}_\tau^s, \widehat{\Sigma}_u^s\right)$, where $\hat{\theta}_\tau^s$ is the estimated coefficient and $\widehat{\Sigma}_u^s$ is the estimated variance-covariance matrix of the local projection residuals u_{t+s}^s, estimated by Newey and West (1987) with lag order $s-1$ due to the serial correlation in the error term induced by the successive leading of the dependent variable in the s-step ahead direct forecasting regression.

Combining the Uncertainty We merge the two distributions of the impulse response coefficient $\hat{\theta}_{\tau,k}^s$ into one distribution. 68 percent confidence intervals are constructed by looking at the 16th and 84th percentile of that distribution. These intervals are then centered around the point estimate $\hat{\theta}_\tau^s$ obtained with the original sample.

D Gertler and Karadi (2015) Monetary Policy Shocks

We construct the monetary policy shocks implied by the proxy VAR used in Gertler and Karadi (2015) using the following procedure.

First, we update the data used in the Gertler and Karadi (2015) baseline VAR. They use monthly data from 1979M7 to 2012M6. We update all time-series to 2019M12. The VAR includes (the log of) industrial production, (the log of) the consumer price index, the one-year government bond rate, and the excess bond premium. As instrument for the monetary policy shock in the VAR we use our updated Miranda-Agrippino and Ricco (2020) shocks series.

Then, we estimate the proxy VAR and compute the implied structural monetary policy shocks, see the appendix of Mertens and Ravn (2013) for details. The identification of monetary policy shocks is achieved by relying on the correlation between the reduced form residuals of the one-year government bond rate and the instrument and that the later is orthogonal to other structural shocks.

E Robustness

One-Step Approach We report the results from the one-step approach to compute the impulse responses of the one-year-ahead GDP growth distribution in Figures E-1 and E-2 for the monetary policy shock and credit spread shock, respectively. In this one-step specification, we run a following one-step quantile local-projection that delivers:

$$q_\tau\left(\bar{Y}_{t+s}\right) = \delta^s_\tau + \theta^s_\tau \ shock_t + \Psi(L)^s_\tau \ controls_t, \quad s = \{0, \dots, S\} \tag{E-1}$$

where $\bar{Y} \equiv \bar{\Delta} y_{t+1,t+12}$. For the monetary policy shock, we use twelve lags of the shock as controls, whereas for the credit spread shock we use one lag of $\left[\ EBP_t,\ \pi^{cpi}_t,\ \Delta y_t,\ i_t\right]'$ as controls to allow for enough observations.

Horizon - Average GDP Growth over the Next Six Months We explore the sensitivity of our results to the choice of a shorter horizon for the growth outlook. In particular, we replicate our results for the choice of average GDP growth over the next six months as a dependent variable in the quantile regression. Figure E-3.

Starting the Local Projection for Credit Spread Shocks in 1986 In Figure E-4, we show results for the two-step and one-step approaches in a specification where credit spread shocks start in 1986 as the monetary policy shock. We will focus on this specification in which the two shocks share the same sample for the following robustness exercises.

Starting the Quantile Regression in 1986 For the IRFs in Figure E-5 the quantile regression starts in 1986 as the local projection..

Industrial Production Instead of Monthly GDP In Figure E-6 we explore the sensitivity of our baseline results to the choice of the dependent variable. In particular, we replace monthly GDP growth by the (month-over-month) growth in industrial production.

Pre-Great-Recession Sample Local projection stops in 2007 September/Q3. To allow for enough observations we do not control for the lags of the shocks. See Figure E-7.

NFCI Instead of the Financial Factor This specification uses the NFCI instead of the financial factor. See Figure E-8.

NFCI Instead of the Financial Factor with Monetary Policy Shocks Available from 1973 This specification is the same as the previous, but uses shocks available from 1973 so that also the local projection starts in 1973. For this exercise, we explore two types of monetary policy shocks. First, we use the Romer and Romer (2004) (RR henceforth) narrative-based monetary shocks provided by Ramey (2016). They regress the federal funds target rate on Greenbook forecasts at each FOMC meeting date and use the residuals as the monetary policy shock. As a second measure, we use the monetary policy shocks identified by Antolín-Díaz and Rubio-Ramírez (2018) (AR henceforth) who add narrative sign restrictions to the VAR model in Uhlig (2005). The sample period runs from January 1973 to December 2007 (RR) and to July 2007 (NAR). See Figure E-9.

Quarterly Specification with Shocks Available from 1973 This specification is the same as the baseline, but at quarterly frequency and thus using official GDP growth from the national accounts to construct one-year-ahead GDP growth. Here the sample period is 1973Q1 to 2007-Q4 (RR) and 2007-Q3 (AR) and 2019-Q4 (Credit Spread Shocks).[20] See Figure E-10.

[20]Both the Romer and Romer (2004) and Antolín-Díaz and Rubio-Ramírez (2018) monetary policy shocks are aggregated to quarterly frequency by adding up the monthly values within each quarter.

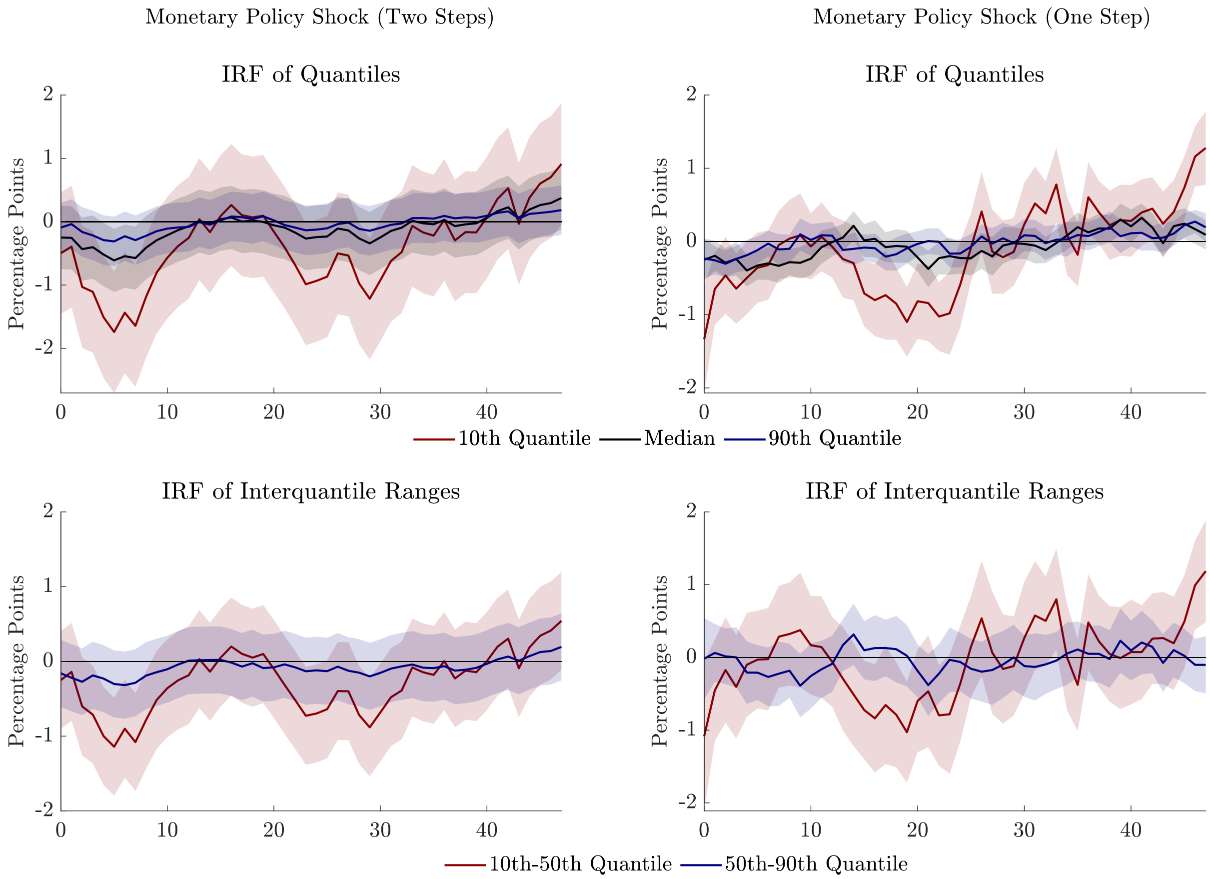

Figure E-1: Impulse Responses of Quantiles of Average GDP Growth over the Next Twelve Months to *Contractionary* Monetary Policy Shocks. Two-Step and One-Step Approach.

Note: Red (dashed) is response of the 10th quantile, black (solid) is the median response, blue (dotted) is response of the 90th quantile. Confidence bands correspond to median response, 68% significance level, based on block-of-blocks bootstrapped standard errors. For the interquantile ranges of the one-step approach, the bands are obtained by summing the uncertainty of the quantile-pair considered (since we do not have an estimate for the covariance) and are thus on the conservative end. The x-axis gives the horizon of the impulse response, in months. The response on the y-axis is measured in percentage points.

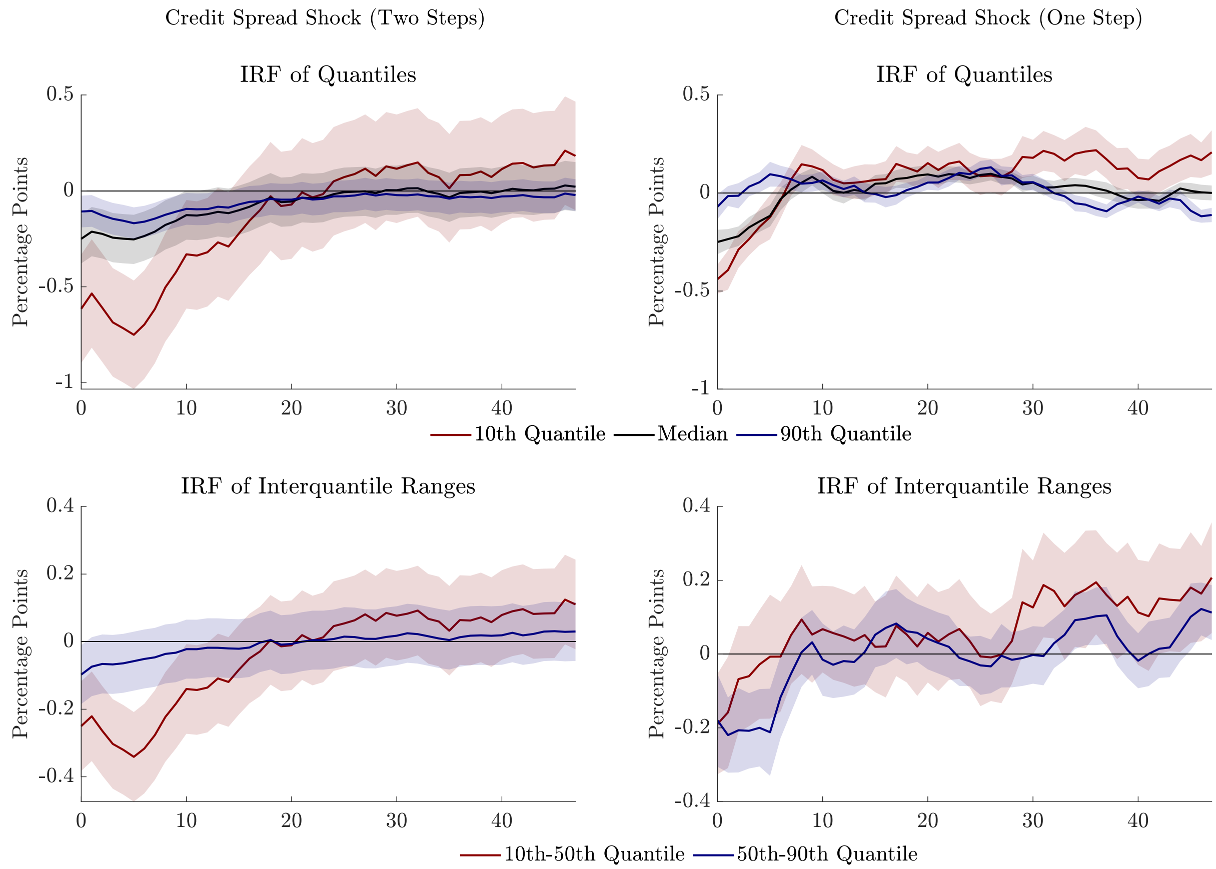

Figure E-2: Impulse Responses of Quantiles of Average GDP Growth over the Next Twelve Months to *Contractionary* Credit Spread Shocks. Two-Step and One-Step Approach.

Note: Red (dashed) is response of the 10th quantile, black (solid) is the median response, blue (dotted) is response of the 90th quantile. Confidence bands correspond to median response, 68% significance level, based on block-of-blocks bootstrapped standard errors. For the interquantile ranges of the one-step approach, the bands are obtained by summing the uncertainty of the quantile-pair considered (since we do not have an estimate for the covariance) and are thus on the conservative end. The x-axis gives the horizon of the impulse response, in months. The response on the y-axis is measured in percentage points.

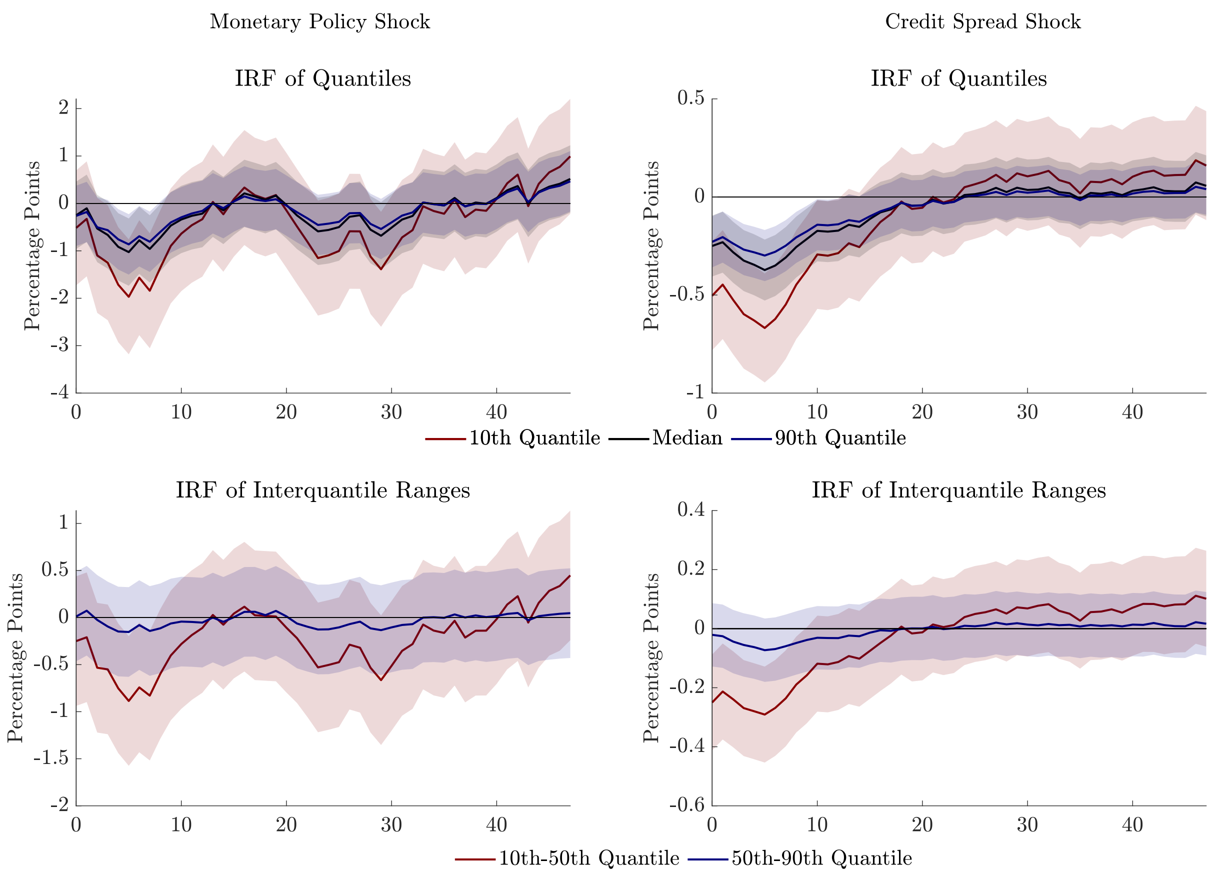

Figure E-3: Impulse Responses of Quantiles of Average GDP Growth over the Next Six Months to *Contractionary* Shocks.

Note: Red (dashed) is response of the 10th quantile, black (solid) is the median response, blue (dotted) is response of the 90th quantile. Confidence bands correspond to median response, 68% significance level, based on Newey-West and block-of-blocks bootstrapped standard errors. The x-axis gives the horizon of the impulse response, in months. The response on the y-axis is measured in percentage points.

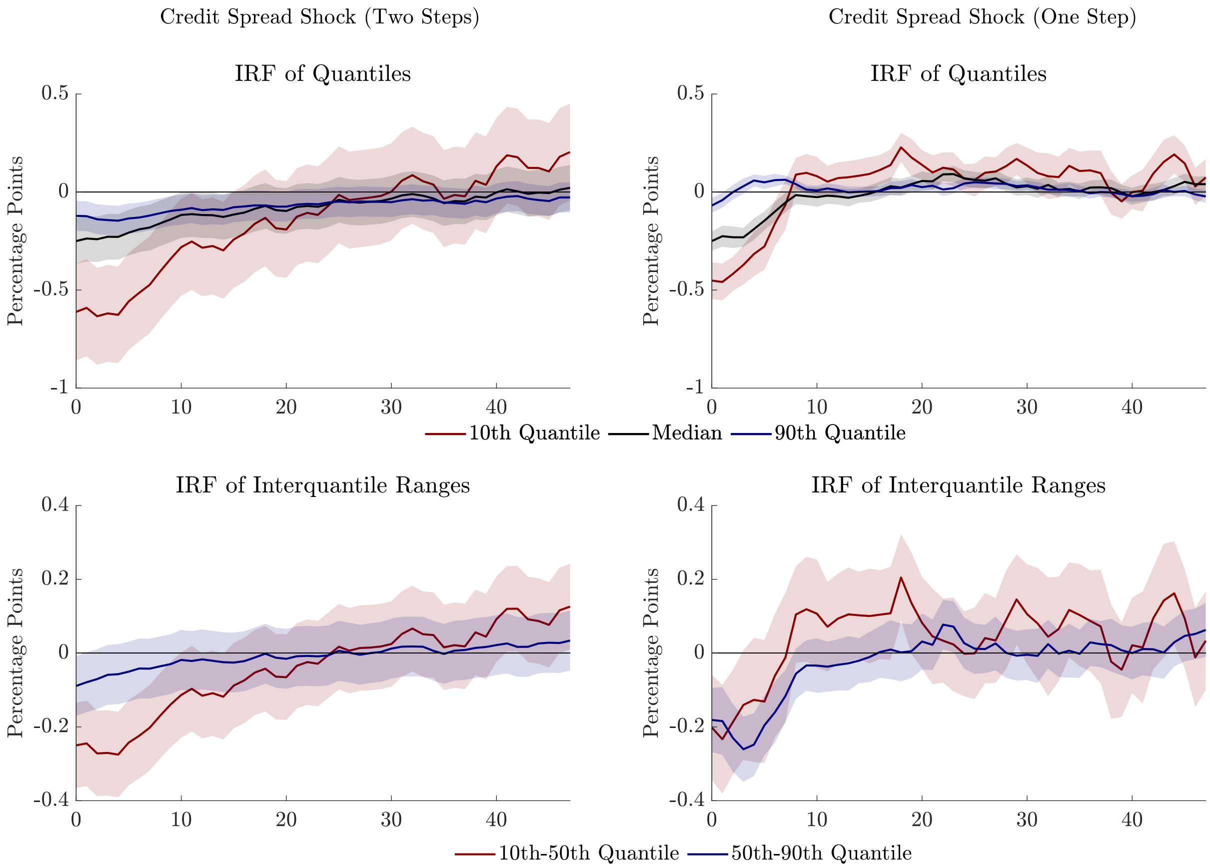

Figure E-4: Impulse Responses of Quantiles of Average GDP Growth over the Next Twelve Months to *Contractionary* Shocks. Local Projection for Credit Spread Shock Starts in 1986. Two-Step and One-Step Approach.

Note: Red (dashed) is response of the 10th quantile, black (solid) is the median response, blue (dotted) is response of the 90th quantile. Confidence bands correspond to median response, 68% significance level, based on Newey-West and block-of-blocks bootstrapped standard errors. The x-axis gives the horizon of the impulse response, in months. The response on the y-axis is measured in percentage points.

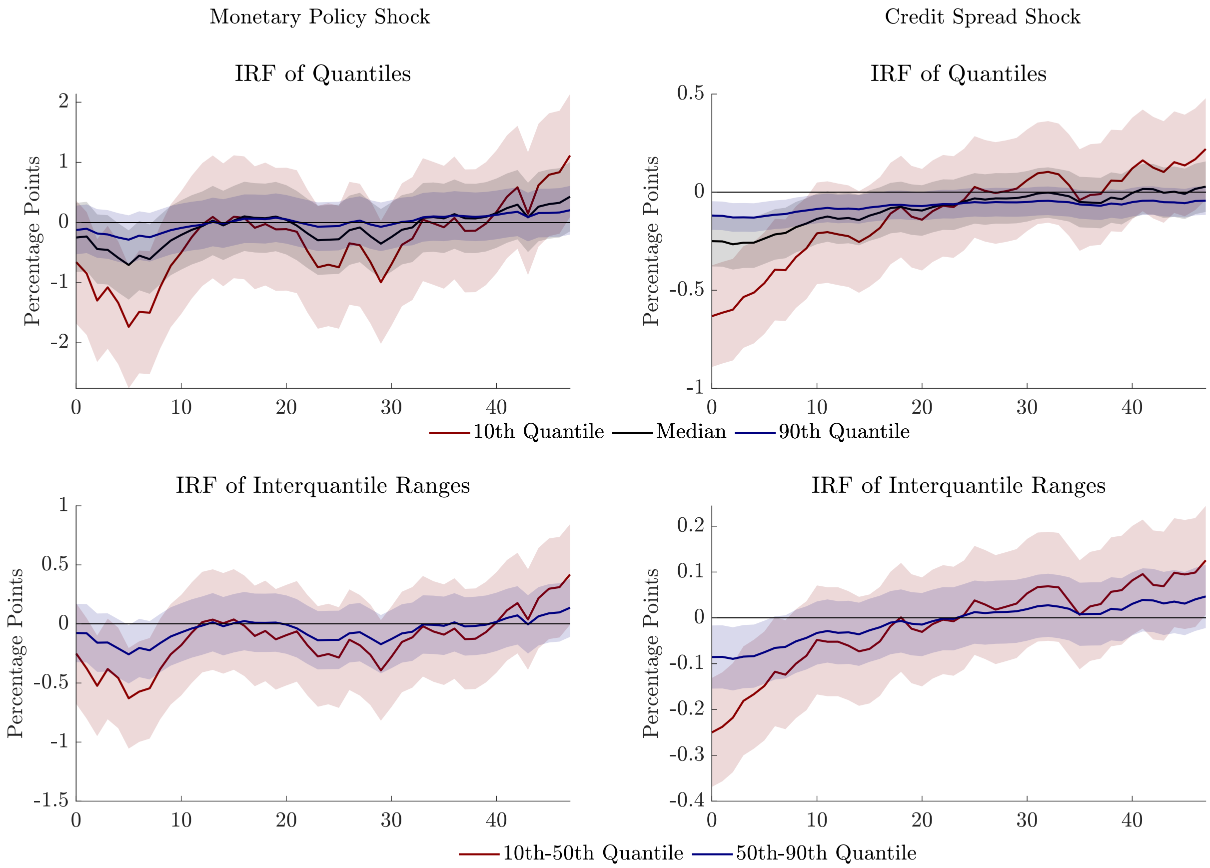

Figure E-5: Impulse Responses of Quantiles of Average GDP Growth over the Next Twelve Months to *Contractionary* Shocks. Estimation of Quantile Regression Starts in 1986.

Note: Red (dashed) is response of the 10th quantile, black (solid) is the median response, blue (dotted) is response of the 90th quantile. Confidence bands correspond to median response, 68% significance level, based on Newey-West and block-of-blocks bootstrapped standard errors. The x-axis gives the horizon of the impulse response, in months. The response on the y-axis is measured in percentage points.

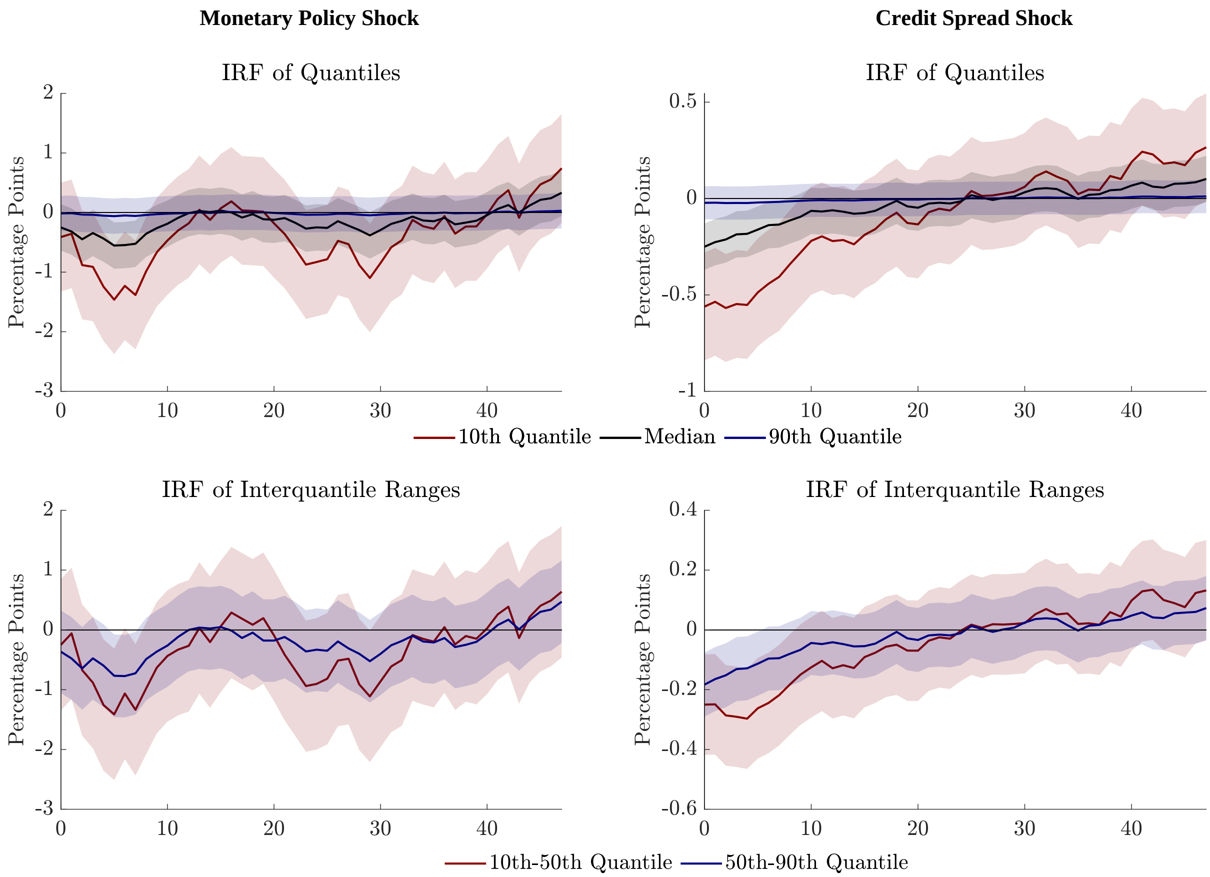

Figure E-6: Impulse Responses of Quantiles of Average Industrial Production Growth over the Next Year to *Contractionary* Shocks. Specification using Industrial Production Instead of Monthly GDP as Dependent Variable.

Note: Red (dashed) is response of the 10th quantile, black (solid) is the median response, blue (dotted) is response of the 90th quantile. Confidence bands correspond to median response, 68% significance level, based on Newey-West and block-of-blocks bootstrapped standard errors. The x-axis gives the horizon of the impulse response, in months. The response on the y-axis is measured in percentage points.

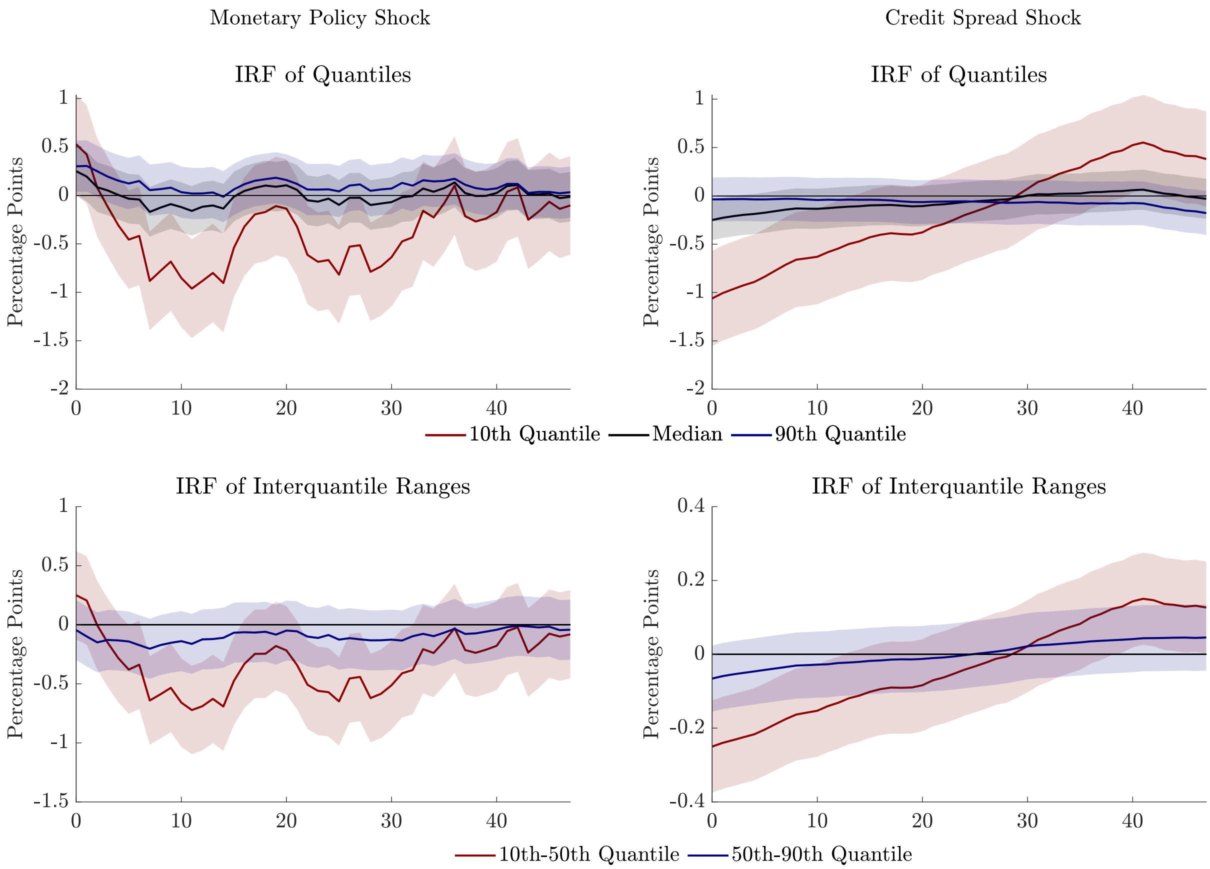

Figure E-7: Impulse Responses of Quantiles of Average GDP Growth over the Next Year to *Contractionary* Shocks. Sample Stopping Before Great Recession.

Note: Red (dashed) is response of the 10th quantile, black (solid) is the median response, blue (dotted) is response of the 90th quantile. Confidence bands correspond to median response, 68% significance level, based on Newey-West and block-of-blocks bootstrapped standard errors. The x-axis gives the horizon of the impulse response, in months. The response on the y-axis is measured in percentage points.

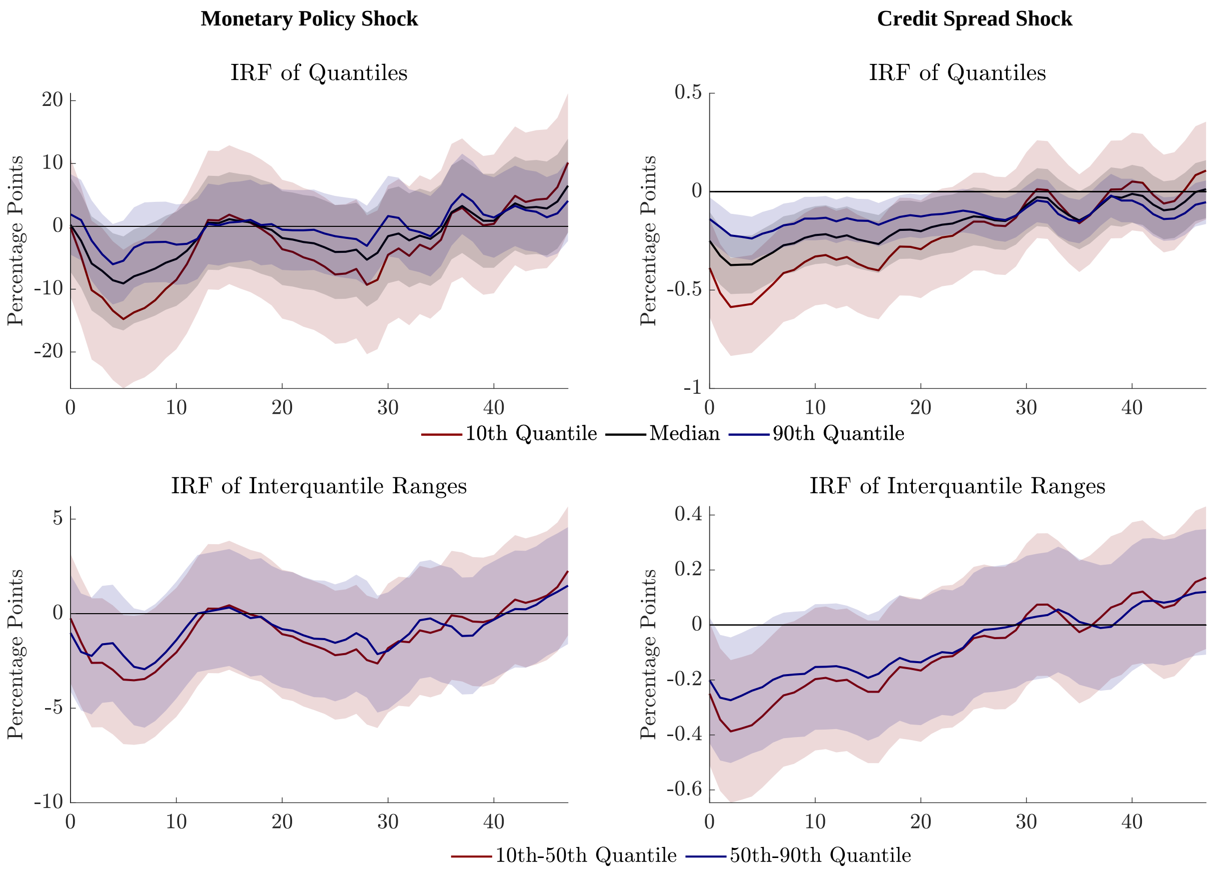

Figure E-8: Impulse Responses of Quantiles of Average GDP Growth over the Next Year to *Contractionary* Shocks. Specification with NFCI Instead of Financial Factor.

Note: Red (dashed) is response of the 10th quantile, black (solid) is the median response, blue (dotted) is response of the 90th quantile. Confidence bands correspond to median response, 68% significance level, based on Newey-West and block-of-blocks bootstrapped standard errors. The x-axis gives the horizon of the impulse response, in months. The response on the y-axis is measured in percentage points.

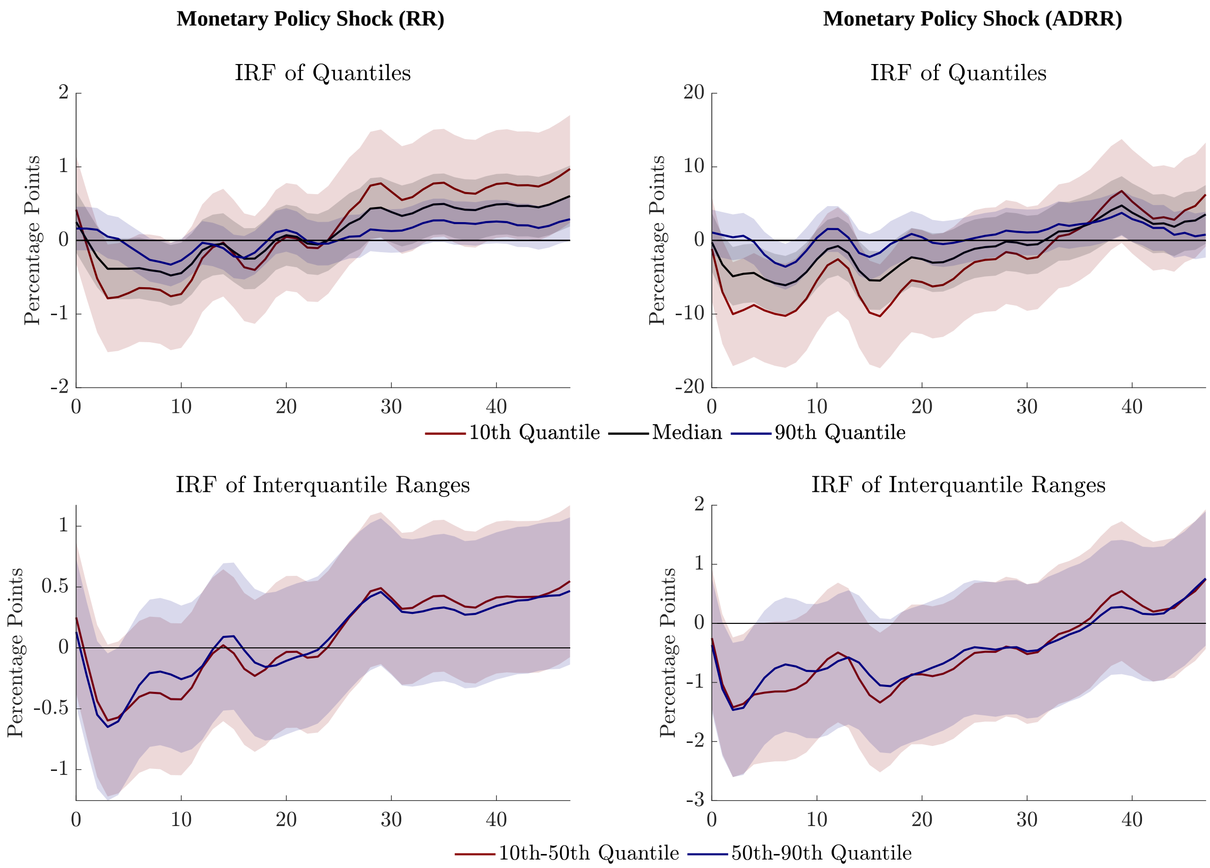

Figure E-9: Impulse Responses of Quantiles of Average GDP Growth over the Next Year to *Contractionary* Shocks. Specification with NFCI Instead of Financial Factor with Monetary Policy Shocks Available from 1973.

Note: Red (dashed) is response of the 10th quantile, black (solid) is the median response, blue (dotted) is response of the 90th quantile. Confidence bands correspond to median response, 68% significance level, based on Newey-West and block-of-blocks bootstrapped standard errors. The x-axis gives the horizon of the impulse response, in months. The response on the y-axis is measured in percentage points.

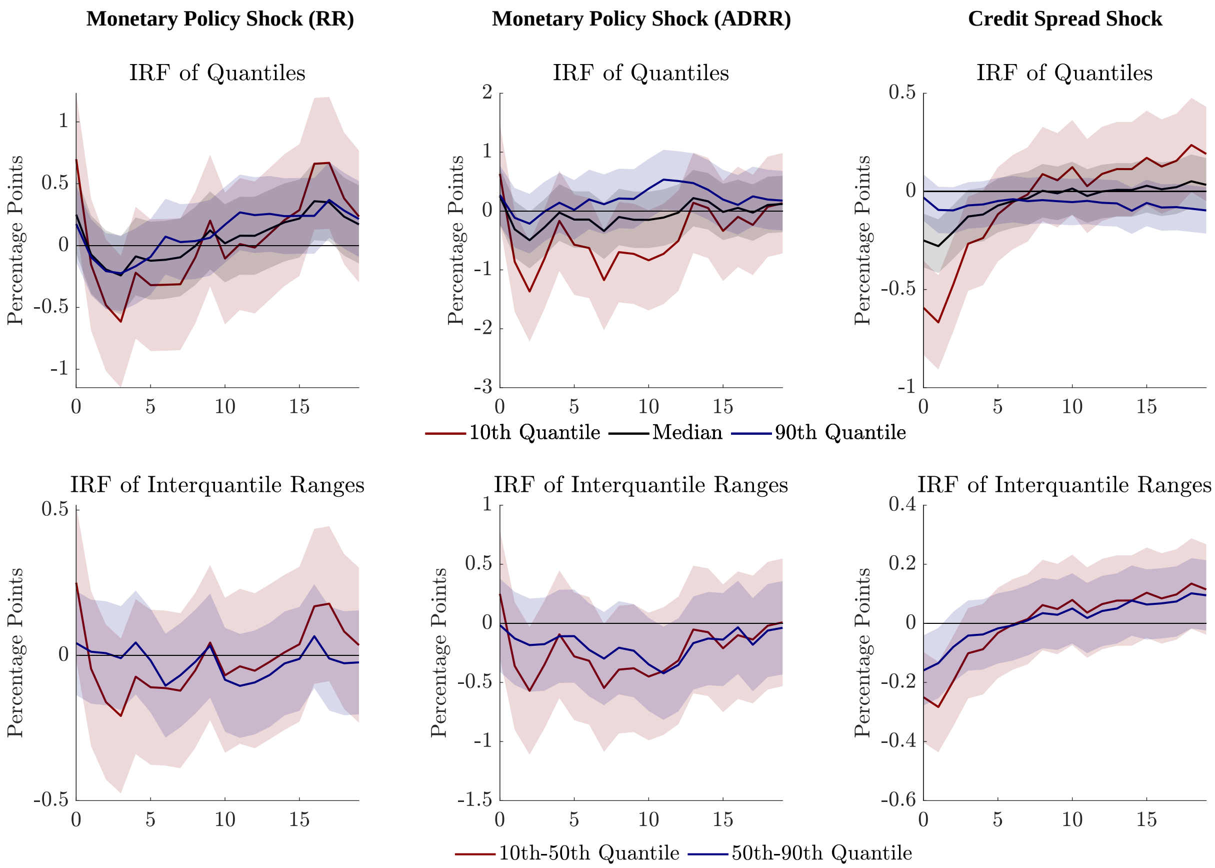

Figure E-10: Impulse Responses of Quantiles of Average GDP Growth over the Next Year to *Contractionary* Shocks. Quarterly Specification (with BEA GDP Growth) with Shocks Starting in 1973:Q1.

Note: Red (dashed) is response of the 10th quantile, black (solid) is the median response, blue (dotted) is response of the 90th quantile. Confidence bands correspond to median response, 68% significance level, based on Newey-West and block-of-blocks bootstrapped standard errors. The x-axis gives the horizon of the impulse response, in months. The response on the y-axis is measured in percentage points.

F Details on Threshold VAR Model

In this section, we provide details on the parameter values of the models presented in Subsection 4.1.

Parameterization The parameters f^* and m^* are chosen such that, on average, the constraint binds in 10% of the sample (on average, across simulations), which is about the percentage of NBER-dated recessions over the January 1973 to May 2020 sample. $\beta_1 = -0.5$ and $\beta_2 = 0.5$. Negative values for β_1 mean that an increase in f_t is associated with a tightening in financial conditions, which depresses growth. Conversely, positive values for β_2 mean that an increase in m_t is associated with an improvement in macroeconomic activity, which fosters growth.

We assume $\beta_0 = 1$, $\alpha_1 = \gamma_1 = 0.8$ for equal persistence of f_t and m_t, and $\alpha_2 = \gamma_2 = -0.5$ as the financial and macroeconomic factor are negatively related, $\alpha_3 = 1$ and $\gamma_3 = -1$ as a positive (contractionary) shock increases the financial factor (tightens financial conditions) and decreases the macroeconomic factor (weakens economic activity).

As to the switching parameters, in model I we assume the following:

$$\alpha_3(f_t, m_t) = \begin{cases} 5, & \text{if } f_t > f^* \;\&\; m_t < m^* \\ 1, & \text{normal state} \end{cases}, \quad \gamma_3(f_t, m_t) = \begin{cases} -5, & \text{if } f_t > f^* \;\&\; m_t < m^* \\ -1, & \text{normal state} \end{cases}$$

In the symmetric model α_3 and γ_3 take on the values of the normal state.